Scenario generation for stochastic programs with tail risk measures

Jamie Fairbrother*, Amanda Turner*, and Stein W. Wallace**

*STOR-i Centre for Doctoral Training, Lancaster University. United Kingdom
**Department of Business and Management Science, Norwegian School of Economics. Norway

November 11, 2015

arXiv:1511.03074v1 [math.OC] 10 Nov 2015

Abstract

Tail risk measures such as Value-at-Risk and Conditional Value-at-Risk are used in stochastic programming to mitigate or reduce the probability of large losses. However, because tail risk measures only depend on the upper tail of a distribution, scenario generation for these problems is difficult as standard methods such as sampling will typically inadequately represent these areas. We present a problem-based approach to scenario generation for stochastic programs which use tail risk measures, and demonstrate this approach on a class of portfolio selection problems.

1 Introduction

Stochastic programming is a tool for making decisions under uncertainty. Stochastic programs are used to model situations where an initial decision must be taken with some information unknown until after the decision has been made. For example, one may want to know how much to invest in a new production technology without knowing exactly the future demand for the product. In stochastic programming, uncertain parameters are modeled as random variables, and one attempts to minimize the expectation or risk measure of some loss function which depends on the initial decision. However, what distinguishes stochastic programming from other stochastic modeling approaches is the ability to explicitly model future decisions based on outcomes of stochastic parameters and initial decisions, and the associated costs of these future decisions. In our example, given an investment decision and a demand, we could model how to distribute this product and the costs of this distribution. The power and flexibility of the stochastic programming approach comes at a price: stochastic programs are usually analytically intractable, and not susceptible to deterministic optimization techniques. See [KW12] for a guide to how stochastic programs are used to model real problems, and [KW94], [BL97] for more general overviews of the subject.

Typically, a stochastic program can only be solved when it is *scenario-based*, that is when the random variables of the problem have finite discrete distributions. For example, stochastic linear programs just become linear programs when the underlying random variables are discrete. In the stochastic programming literature, the mass points of these random variables are referred to as *scenarios*, the discrete distribution as the *scenario set* and the construction of this as *scenario generation*. Scenario generation can consist of discretizing a continuous probability distribution, or directly modeling the uncertain quantities as discrete random variables. The more scenarios in a set, the more computational power that is required to solve the problem. The key issue of scenario generation is how to represent the uncertainty to ensure that the solution to the problem is reliable, while keeping the number of scenarios low so that the problem is computational tractable. See [KW07] for methods of evaluating scenario generation methods and a discussion of what constitutes a reliable solution.

Minimizing the expectation of a loss function can be thought of as minimizing the long-term costs of a system. This is appropriate when the initial decision is going to be used again and again, and large losses do not matter in the short term. For example, a news vendor may have to decide on a daily order of items to which they are committed for some period of time. In other cases, the decision may be only used a few times, and the occurrence of large losses may be lead to bankruptcy. In this latter case, minimizing the expectation alone is not appropriate as this does not necessarily mitigate against large losses. The usual action of recourse in this case is to use some sort of *risk measure* which quantifies in some way the likelihood and severity of potential large losses. In these problems we try to find a decision which appropriately balances in the expectation against risk.

In this paper we are interested in problems which use *tail risk measures*. A precise definition of a tail-risk measure will be given in Section 2 but for now, one can think of a tail risk measure as a function of a random variable which only depends on the upper tail of its distribution function. Examples of tail risk measure include the Value-at-Risk [Jor96] and the Conditional Value-at-Risk [RU00], both of which are commonly used in financial contexts. The problem of scenario generation is particularly acute when the scenarios are being used to calculate the value of a tail risk measure. This is because standard scenario generation methods will not produce many scenarios in the tail of the loss function and so it is inadequately represented.

The most basic approach to discretization is to simply use a random sample from the true distribution. This has desirable asymptotic properties [KR93], [Sha03], but may require large sample sizes to ensure the reliability of the solutions it yields. This can be mitigated somewhat by using variance reduction techniques such as stratified sampling and importance sampling [LSW06]. Sampling also has the advantage that it can be used to construct confidence intervals on the true solution value [MMW99]. Another approach to discretization is to construct a distribution whose distance from the true distribution, with respect to some probability metric, is small [Pfl01], [HR09]. These approaches tend to yield better and much more stable solutions to stochastic programs than does sampling.

A characteristic of both of these approaches to scenario generation is that they are *distribution-based*; that is, they only aim to approximate a distribution and are divorced from the stochastic program for which they are producing scenarios. By exploiting the structure of a problem, it may be possible to find a more parsimonious representation of the uncertainty. Note that such a *problem-based* approach may not yield a discrete distribution which is close to the true distribution in a probabilistic sense; the aim is only to find a discrete distribution which yields a high quality solution to our problem.

A set of approaches which move away from the purely distribution-based paradigm of scenario generation are *constructive methods*. In these approaches, the modeler does not use a full probability distribution for the uncertain problem parameters but specifies a set of target statistical properties they believe the distribution satisfies, and generates a scenario set with these target properties. This approach was first proposed in [HW01], where it is postulated that the solution to a stochastic program will depend largely on a small set of statistical properties of the random variables, specific to that problem. That is, if we can generate a scenario set with the required properties, this should yield good solutions in our stochastic program even if the true distribution is significantly different. For example, it is known that for the classical Markowitz problem [Mar52] the first two moments of the return distributions determine exactly the solution. Constructive approaches have gained much popularity because they simplify the stochastic modeling of the uncertain parameters. In particular they eliminate the need to fit parametric stochastic models. Other constructive approaches can be found in [HKW03], and [KW11]. However, the major draw-back with constructive approaches is that it is not always clear which properties are important for a given problem. Finding out which properties are important is therefore an important part of the analysis.

In this paper, we present a general problem-based approach to scenario generation for stochastic programs which use tail risk measures. We observe that the value of any tail risk measure depends only on scenarios confined to an area that we call the *risk region*. This means that all scenarios not in the risk region can be aggregated into a single point. By concentrating almost all scenarios in the risk region, we can calculate the value a tail risk measure more accurately. One feature of the risk region is that the more constrained our problem, the smaller it becomes, and so the more useful our methodology. However, finding the risk region is difficult as it is determined by both the problem and the distribution

of the uncertain parameters.

We demonstrate our methodology for portfolio selection problems where the assets are assumed to have returns which are elliptically distributed. For this type of problem we are able to characterize the risk region in a convenient way. We will show that the risk region depends only on the conic hull of our feasible region. Another useful property of the portfolio selection problem is the linearity (affinity) of the loss function. This means that all scenarios not in the risk region can be aggregated while preserving the overall expected return.

Some ideas in this paper are similar to those in [GBM12]. In that paper, the authors, like us, observe that only scenarios which have a loss in the tail of the distribution are used in the calculation of the tail risk measure. However, while we use this observation to construct a scenario set, they exploit this property to solve a problem which uses the β-CVaR risk measure for a given scenario set. Their approach is to iteratively solve the problem with a subset of scenarios, identify the scenarios which have loss in the tail, update their scenario set appropriately and resolve, until the true solution has been found.

This paper is organized as follows: in Section 2 we define tail risk measures and their associated risk regions; in Section 3 we discuss how these risk regions can be exploited for the purposes of scenario generation and scenario reduction; in Section 4 we prove that our scenario generation method is consistent with sampling, in Section 5 and Section 6 we provide a proof of concept for our methodology: we give convenient characterizations for risk regions for a class of portfolio selection problems and present numerical tests which compare our methodology against basic sampling; finally in Section 7 we summarize our results make some concluding remarks.

2 Tail risk measures

2.1 Tail risk of random variables

Suppose that we have an uncertain quantity representing some loss, and we would like to somehow quantify the riskiness of this quantity. We model the uncertain quantity as a random variable and take a risk measure to be any function of a random variable. The following definition is taken from [Tas02].

Definition 2.1 (Risk Measure)**.** Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, and V be a non-empty set of $\mathcal{F}$-measurable real-valued random variables[1]. Then, a risk measure is some function $\rho : V \to \mathbb{R} \cup \{\infty\}$.

However, for a risk measures to be useful, it should in some way penalize potential large losses. For example, in the classical Markowitz problem [Mar52], the uncertain quantity is the return of a portfolio of financial assets, and the measure of risk is the variance of that return. By choosing a portfolio with a low variance, we reduce the probability of larges losses as a direct consequence of Chebyshev's inequality (see for instance [Bil95]). Various criteria for risk measures have been proposed; in [ADEH99] a *coherent risk measure* is defined to be a risk measure which satisfies axioms such as positive homogeneity and subadditivity; another perhaps desirable criterion for risk measures is that the risk measure is consistent with respect to first and second order stochastic dominance, see [OR02] for instance.

Besides not satisfying some of the above criteria, a major drawback with using variance as a measure is that it penalizes all large deviations from the mean, that is, it penalizes large profits as well as large losses. This problem can be overcome by using *downside* risk measures such as the semi-variance, which only penalize losses above the mean. However, if we are truly interested in rare or extreme losses, using a risk measure which still depends on the main body of the distribution such as semi-variance may give us distorted or over-optimistic results.

These considerations motivate the idea of using risk measures which depend only on the upper tail of the distribution. To be more precise, the upper tail of a distribution consists of outcomes with a loss greater than or equal to some quantile of the underlying distribution function.

Definition 2.2 (Quantile Function)**.** Suppose Z is a random variable with distribution function F_Z.

[1] We implicitly assume throughout that V is large enough to contain all constructed random variables

Then the generalized inverse distribution function, or *quantile function* is defined as follows:

$$F_Z^{-1} : (0,1] \to \mathbb{R} \cup \{\infty\}$$
$$\beta \mapsto \inf\{x \in \mathbb{R} : F_Z(x) \geq \beta\}.$$

Definition 2.3 (Tail Risk Measure)**.** Let $\rho_\beta : V \to \mathbb{R} \cup \{\infty\}$ be a risk measure as above, then ρ_β is a β-tail risk measure if $\rho_\beta(Z)$ depends only on the restriction of quantile function of Z above β, that is $F_Z^{-1} \restriction_{[\beta,1]}$.

To show that ρ_β is a β-tail risk measure, we must show that $\rho_\beta(Z)$ can be written as a function of the quantile function above or equal to β. Two very popular tail risk measures are the value-at-risk [Jor96] and the conditional value-at-risk [RU02]:

Example 2.4 (Value at risk)**.** Let Z be a random variable, and $0 < \beta < 1$. Then, the $\beta-$VaR for Z is defined to be the β-quantile of Z:

$$\beta\text{-VaR}(Z) := F_Z^{-1}(\beta).$$

Example 2.5 (Conditional value at risk)**.** Let Z be a random variable, and $0 < \beta < 1$. Then, the β-CVaR can be thought roughly as the conditional expectation of a random variable above its β-quantile. The following alternative characterization of β-CVaR [AT02] shows directly that it is a β-tail risk measure.

$$\beta\text{-CVaR}(Z) = \int_\beta^1 F_Z^{-1}(u)\, du.$$

The observation that we exploit for this work is that very different random variables will have the same β-tail risk measure as long as their β-tails are the same. Such a situation is illustrated in Figure 1 for two discrete random variables.

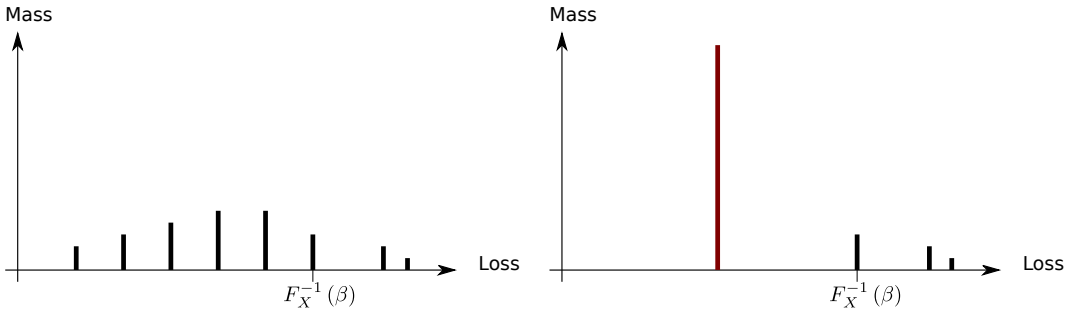

Figure 1: Two very different random variables with identical β-tails

When showing that two distributions have the same β-tails, it is convenient to use distribution functions rather than quantile functions. An equivalent condition for showing that two random variables Z_1 and Z_2 have the same β-tail, that is $F_{Z_1}^{-1}(u) = F_{Z_2}^{-1}(u)$ for all $\beta \leq u \leq 1$, is the following:

$$F_{Z_1}^{-1}(\beta) = F_{Z_2}^{-1}(\beta) \text{ and } F_{Z_1}(z) = F_{Z_2}(z) \text{ for all } z \geq F_{Z_1}^{-1}(\beta).$$

2.2 Optimization and risk regions

In the optimization context, we suppose that the loss depends on some decision $x \in \mathcal{X} \subseteq \mathbb{R}^k$ and the outcome of some latent random vector Y with support $\mathcal{Y} \subseteq \mathbb{R}^d$, defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and which is independent of x. That is, we suppose our loss is determined by some function, $f : \mathcal{X} \times \mathbb{R}^d \to \mathbb{R}$, which we refer to as the *loss function*. For a given decision $x \in \mathcal{X}$, the random variable associated with the loss is thus $f(x, Y)$.

We are typically interested in optimization problems which attempt balance expected loss or profit against the risk of a decision. For example, see the problem statements in Section 5.1. To avoid repeated

use of cumbersome notation we introduce the following short-hand for distribution and quantile functions:

$$F_x(z) := F_{f(x,Y)}(z) = \mathbb{P}\left(f(x,Y) \leq z\right),$$
$$F_x^{-1}(\beta) := F^{-1}_{f(x,Y)}(\beta) = \inf\{z \in \mathbb{R} : \ F_x(z) \geq \beta\}.$$

Since tail risk measures depend only on those outcomes which are in the β-tail, we aim to identify the region of the support which lead to a loss in the β-tails for some decision.

Definition 2.6 (Risk region)**.** For $0 < \beta < 1$ the β-risk region with respect to the decision $x \in \mathcal{X}$ is defined as follows:

$$\mathcal{R}_x(\beta) = \{y \in \mathbb{R}^d : F_x\left(f(x,y)\right) \geq \beta\},$$

or equivalently

$$\mathcal{R}_x(\beta) = \{y \in \mathbb{R}^d : f(x,y) \geq F_x^{-1}(\beta)\}. \tag{1}$$

The risk region with respect to the feasible region $\mathcal{X} \subset \mathbb{R}^k$ is defined to be:

$$\mathcal{R}_{\mathcal{X}}(\beta) = \bigcup_{x \in \mathcal{X}} \mathcal{R}_x(\beta).$$

The complement of this region is called the *non-risk region*. This can also be written

$$\mathcal{R}_{\mathcal{X}}(\beta)^c = \bigcap_{x \in \mathcal{X}} \mathcal{R}_x(\beta)^c. \tag{2}$$

The following basic properties of the risk region follow directly from the definition.

$$\text{(i) } 0 < \beta' < \beta < 1 \ \Rightarrow \ \mathcal{R}_{\mathcal{X}}(\beta) \subseteq \mathcal{R}_{\mathcal{X}}(\beta'); \tag{3}$$
$$\text{(ii) } \mathcal{X}' \subset \mathcal{X} \ \Rightarrow \mathcal{R}_{\mathcal{X}'}(\beta) \subseteq \mathcal{R}_{\mathcal{X}}(\beta); \tag{4}$$
$$\text{(iii) If } y \mapsto f(x,y) \text{ is continuous then } \mathcal{R}_x(\beta) \text{ is closed and } \mathcal{R}_x(\beta)^c \text{ is open.} \tag{5}$$

We now state a technical property and prove that this ensures the distribution of the random vector in a given region completely determines the value of a tail risk measure. We then identify conditions which guarantee that this risk region satisfies this property.

Definition 2.7 (Aggregation condition)**.** Suppose that $\mathcal{R}_{\mathcal{X}}(\beta) \subseteq \mathcal{R} \subset \mathbb{R}^d$ and that for all $x \in \mathcal{X}$, $\mathcal{R}$ satisfies the following condition:

$$\mathbb{P}\left(Y \in \{y : z' < f(x,y) \leq F_x^{-1}(\beta)\} \cap \mathcal{R}\right) > 0 \qquad \forall \ z' < F_x^{-1}(\beta). \tag{6}$$

Then $\mathcal{R}$ is said to satisfy the *β-aggregation condition.*

Theorem 2.8. Suppose that $\mathcal{R}_{\mathcal{X}} \subseteq \mathcal{R} \subset \mathbb{R}^d$ satisfies the β-aggregation condition and that $\tilde{Y}$ is a random vector for which

$$\mathbb{P}\left(Y \in \mathcal{A}\right) = \mathbb{P}\left(\tilde{Y} \in \mathcal{A}\right) \qquad \text{for any } \mathcal{A} \subseteq \mathcal{R}. \tag{7}$$

Then $\rho_\beta\left((f(x,Y)\right) = \rho_\beta\left(f(x,\tilde{Y})\right)$ for all $x \in \mathcal{X}$, for any β-tail risk measure ρ_β.

Proof. Fix $x \in \mathcal{X}$. To show that $\rho_\beta\left(f(x,Y)\right) = \rho_\beta\left(f(x,\tilde{Y})\right)$ we must show that the β-quantile and the β-tail distributions of $f(x,Y)$ and $f(x,\tilde{Y})$ are the same. The following two conditions are necessary and sufficient for this to occur:

$$F_x(z) = F_{f(x,\tilde{Y})}(z) \qquad \forall \ z \geq F_x^{-1}(\beta),$$
$$F_{f(x,\tilde{Y})}(z) < \beta \qquad \forall z < F_x^{-1}(\beta).$$

Suppose $z' \geq F_x^{-1}(\beta)$. Then,

$$\begin{aligned}
F_{f(x,\tilde{Y})}(z') &= \mathbb{P}\left(\tilde{Y} \in \{y:\ f(x,y) \leq z'\}\right) \\
&= \mathbb{P}\left(\tilde{Y} \in \underbrace{\mathcal{R}^c \cap \{y:\ f(x,y) \leq z'\}}_{=\mathcal{R}^c}\right) + \mathbb{P}\left(\tilde{Y} \in \underbrace{\mathcal{R} \cap \{y:\ f(x,Y) \leq z'\}}_{\subset \mathcal{R}}\right) \\
&= \mathbb{P}\left(Y \in \mathcal{R}^c\right) + \mathbb{P}\left(Y \in \mathcal{R} \cap \{y : f(x,y) \leq z'\}\right) \\
&= \mathbb{P}\left(Y \in \{y : f(x,y) \leq z'\}\right) \\
&= F_x(z')
\end{aligned}$$

as required.

Now suppose $z' < F_x^{-1}(\beta)$. There are two cases; in the first instance suppose $\mathbb{P}\left(f(x,Y) = F_x^{-1}(\beta)\right) > 0$, then we have:

$$\begin{aligned}
F_{f(x,\tilde{Y})}(z') &\leq \mathbb{P}\left(f(x,\tilde{Y}) < F_x^{-1}(\beta)\right) \\
&= \mathbb{P}\left(f(x,Y) < F_x^{-1}(\beta)\right) \\
&< \beta,
\end{aligned}$$

as required. In the case where $\mathbb{P}\left(f(x,Y) = F_x^{-1}(\beta)\right) = 0$ we have:

$$\begin{aligned}
F_{f(x,\tilde{Y})}(z') &= \mathbb{P}\left(\tilde{Y} \in \{y : f(x,y) \leq z'\}\right) \\
&\leq \mathbb{P}\left(\tilde{Y} \in \mathcal{R}^c \cup \{y : f(x,y) \leq z'\}\right) \\
&= \mathbb{P}\left(f(x,\tilde{Y}) \leq F_x^{-1}(\beta)\right) - \mathbb{P}\left(\tilde{Y} \in \underbrace{\mathcal{R} \cap \{y : z' < f(x,y) \leq F_x^{-1}(\beta)\}}_{\subseteq \mathcal{R}}\right) \\
&= 1 - \mathbb{P}\left(f(x,\tilde{Y}) > F_x^{-1}(\beta)\right) - \underbrace{\mathbb{P}\left(Y \in \mathcal{R} \cap \{y : z' < f(x,y) \leq F_x^{-1}(\beta)\}\right)}_{>0 \text{ by (6)}} \\
&< 1 - \underbrace{\mathbb{P}\left(f(x,Y) > F_x^{-1}(\beta)\right)}_{=1-\beta} \\
&= \beta,
\end{aligned}$$

as required. □

The motivation for the term *aggregation condition* follows from the above result: if a set satisfies the aggregation condition then we can transform the probability distribution of Y so that all the mass in the complement of this set is aggregated into a single point without affecting the value of the tail risk measure. This property is particularly relevant to scenario generation as if we have such a set, then all points in its complement can be aggregated, reducing the size of the stochastic program.

The β-aggregation condition is difficult to verify directly. The following shows that it immediately holds for $\mathcal{R}_{\mathcal{X}}(\beta')$ when $\beta' < \beta$.

Proposition 2.9. Suppose $\beta' < \beta$. Then, $\mathcal{R}_{\mathcal{X}}(\beta')$ satisfies the β-aggregation condition. That is for all $x \in \mathcal{X}$

$$\mathbb{P}\left(Y \in \{y : z' \leq f(x,y) \leq F_x^{-1}(\beta)\} \cap \mathcal{R}_{\mathcal{X}}(\beta')\right) > 0 \qquad \forall\ z' < F_x^{-1}(\beta).$$

Proof. Fix $x \in \mathcal{X}$.

Case 1: $F_x^{-1}(\beta') = F_x^{-1}(\beta)$.

In this case, the distribution function F_x has a discontinuity at $z = F_x^{-1}(\beta)$, that is $\mathbb{P}(f(x,Y) = z) > 0$.

Therefore, for $z' < z$ we have

$$\begin{aligned}\mathbb{P}\left(Y \in \{y : z' \leq f(x,y) \leq F_x^{-1}(\beta)\} \cap \mathcal{R}_{\mathcal{X}}(\beta')\right) &\geq \mathbb{P}\left(f(x,Y) = z\right) \\ &> 0\end{aligned}$$

as required.

Case 2: $F_x^{-1}(\beta') < F_x^{-1}(\beta)$.

In this case for all $F_x^{-1}(\beta') < z' < F_x^{-1}(\beta)$, we have $\{y : z' < f(x,y) \leq F_x^{-1}(\beta)\} \subset \mathcal{R}_{\mathcal{X}}(\beta')$ and so

$$\begin{aligned}\mathbb{P}\left(Y \in \{y : z' \leq f(x,y) \leq F_x^{-1}(\beta)\} \cap \mathcal{R}_{\mathcal{X}}(\beta')\right) &= \mathbb{P}\left(z' \leq f(x,Y) \leq F_x^{-1}(\beta)\right) \\ &> 0.\end{aligned}$$

□

From this point onward, for simplicity, we drop the β from our risk region notation and refer to the aggregation condition rather than the β-aggregation condition.

All sets satisfying the aggregation condition must contain the risk region, however, the aggregation condition does not necessarily hold for the risk region itself. It is guaranteed to hold if Y has a discrete distribution, since in this case for all $x \in \mathcal{X}$ and $z' < F_x^{-1}(\beta)$ we have:

$$\begin{aligned}\mathbb{P}\left(Y \in \{y : z' < f(x,y) \leq F_x^{-1}(\beta)\} \cap \mathcal{R}_{\mathcal{X}}\right) &\geq \mathbb{P}\left(f(x,Y) = F_x^{-1}(\beta)\right) \\ &> 0.\end{aligned}$$

In the non-discrete case we must impose extra conditions on the problem to avoid some degenerate cases. Recall that $\mathcal{Y}$ denotes the support of the random vector Y.

Proposition 2.10. Suppose the following conditions hold:

(i) $\operatorname{int}(\mathcal{Y}) \cap \operatorname{int}(\mathcal{R}_{\mathcal{X}})$ is connected

(ii) $y \mapsto f(x,y)$ is continuous for all $x \in \mathcal{X}$

(iii) For each $x \in \mathcal{X}$ there exists $x' \in \mathcal{X}$ such that

$$\operatorname{int}(\mathcal{Y}) \cap \operatorname{int}(\mathcal{R}_x \cap \mathcal{R}_{x'}) \neq \emptyset \text{ and } \operatorname{int}(\mathcal{Y}) \cap \operatorname{int}(\mathcal{R}_{x'} \setminus \mathcal{R}_x) \neq \emptyset \tag{8}$$

Then the risk region $\mathcal{R}_{\mathcal{X}}$ satisfies the aggregation condition.

Proof. Fix $x \in \mathcal{X}$ and $z' < F_x^{-1}(\beta)$. Pick $x' \in \mathcal{X}$ such that (8) holds. Also, let $y_0 \in \operatorname{int}(\mathcal{Y}) \cap \operatorname{int}(\mathcal{R}_{x'} \setminus \mathcal{R}_x)$ and $y_1 \in \operatorname{int}(\mathcal{Y}) \cap \operatorname{int}(\mathcal{R}_x \cap \mathcal{R}_{x'})$. Since $\operatorname{int}(\mathcal{Y}) \cap \operatorname{int}(\mathcal{R}_{\mathcal{X}})$ is connected there exists continuous $\gamma : [0,1] \to \operatorname{int}(\mathcal{Y}) \cap \operatorname{int}(\mathcal{R}_{\mathcal{X}})$ such that $\gamma(0) = y_0$ and $\gamma(1) = y_1$. Now, $f(x,y_0) < F_x^{-1}(\beta)$ and $f(x,y_1) \geq F_x^{-1}(\beta)$ and so given that $t \mapsto f(x,\gamma(t))$ is continuous there must exist $0 < t < 1$ such that $z' < f(x,\gamma(t)) < F_x^{-1}(\beta)$. That is $\operatorname{int}(\mathcal{Y}) \cap \operatorname{int}(\mathcal{R}_{\mathcal{X}}) \cap \{y : z' < f(x,y) < F_x^{-1}(\beta)\}$ is non-empty. This is a non-empty open set contained in the support of Y and so has positive probability, hence the aggregation condition holds. □

The following Proposition gives a condition under which the non-risk region is convex. This is useful as if we can find some points in the non-risk region, then the complement of the convex hull of these points will contain the risk region.

Proposition 2.11. Suppose that for each $x \in \mathcal{X}$ the function $y \mapsto f(x,y)$ is convex. Then the non-risk region $\mathcal{R}_{\mathcal{X}}^c$ is convex.

Proof. For $x \in \mathcal{X}$, if $y \mapsto f(x,y)$ is convex then the set $\mathcal{R}_x^c = \{y \in \mathbb{R}^d : f(x,y) < F_x^{-1}(\beta)\}$ must be convex. The arbitrary intersection of convex sets is convex, hence $\mathcal{R}_{\mathcal{X}}^c = \bigcap_{x \in \mathcal{X}} \mathcal{R}_x^c$ is convex. □

This convexity condition is held by a large class of stochastic programs, for instance, all two-stage linear recourse problems with fixed recourse will have this property (see, for instance,[BL97]).

The random vector in the following definition plays a special role in our theory.

Definition 2.12 (Aggregated random vector)**.** For some set $\mathcal{R} \subset \mathbb{R}^d$ satisfying the aggregation condition, the *aggregated random vector* is defined as follows:

$$\psi_{\mathcal{R}}(Y) := \begin{cases} Y & \text{if } Y \in \mathcal{R}, \\ \mathbb{E}\left[\, Y | Y \in \mathcal{R}^c \,\right] & \text{otherwise.} \end{cases}$$

If we have $\mathbb{E}\left[Y | Y \in \mathcal{R}^c\right] \in \mathcal{R}^c$ then Theorem 2.8 guarantees that $\rho_\beta\left(f\left(x, \psi_{\mathcal{R}}(Y)\right)\right) = \rho_\beta\left(f\left(x, Y\right)\right)$ for all $x \in \mathcal{X}$. For example, the conditions of Proposition 2.11 will guarantee this. As well as preserving the value of the tail risk measure, the function $\psi_{\mathcal{R}}$ will preserve the expectation for affine cost functions.

Corollary 2.13. Suppose for each $x \in \mathcal{X}$ the function $y \mapsto f(x, y)$ is affine and for a set $\mathcal{R} \subset \mathbb{R}^d$ satisfying the aggregation condition we have that

$$\mathbb{E}\left[Y | Y \in \mathcal{R}^c\right] \in \mathcal{R}^c$$

Then,

$$\rho_\beta\left(f\left(x, \psi_{\mathcal{R}}(Y)\right)\right) = \rho_\beta\left(f\left(x, Y\right)\right), \tag{9}$$
$$\mathbb{E}\left[f\left(x, \psi_{\mathcal{R}^c}\left(Y\right)\right)\right] = \mathbb{E}\left[f(x, Y)\right], \tag{10}$$

for all $x \in \mathcal{X}$.

Proof. The equality (9) follows immediately from Theorem 2.8. For the expectation function we have

$$\begin{aligned} \mathbb{E}\left[\psi_{\mathcal{R}}(Y)\right] &= \mathbb{P}\left(Y \in \mathcal{R}\right) \mathbb{E}\left[\psi_{\mathcal{R}}(Y) | Y \in \mathcal{R}\right] + \mathbb{P}\left(Y \in \mathcal{R}^c\right) \mathbb{E}\left[\psi_{\mathcal{R}}(Y) | Y \in \mathcal{R}^c\right] \\ &= \mathbb{P}\left(Y \in \mathcal{R}\right) \mathbb{E}\left[Y | Y \in \mathcal{R}\right] + \mathbb{P}\left(Y \in \mathcal{R}^c\right) \mathbb{E}\left[Y | Y \in \mathcal{R}^c\right] \\ &= \mathbb{E}\left[Y\right]. \end{aligned}$$

Since $y \mapsto f(x, y)$ is affine this means that

$$\begin{aligned} \mathbb{E}\left[f(x, \psi_{\mathcal{R}}(Y))\right] &= f(x, \mathbb{E}\left[\psi_{\mathcal{R}}(Y)\right] \\ &= f(x, \mathbb{E}\left[Y\right]) \\ &= \mathbb{E}\left[f(x, Y)\right]. \end{aligned}$$

□

3 Scenario generation

In the previous section, we showed that under mild conditions the value of a tail risk measure only depends on the distribution of outcomes in the risk region. In this Section we present and discuss how this feature of tail risk measures may be exploited for the purposes of scenario generation and scenario reduction. We propose and give particular prominence to a scenario generation method which we will call *aggregation sampling*, discuss the efficiency of this algorithm and briefly discuss other possible methods of exploiting risk regions.

Our general approach is as follows: for scenario generation we prioritize the construction of scenarios in the risk region to allow one to better approximate the value of the β-tail risk measure; for scenario reduction we reduce the number of scenarios in the non-risk region which are in some sense redundant for computing the value of the β-tail risk measure. We assume throughout this section that our scenario sets are constructed from some underlying probabilistic model from which we can draw samples, and for which we can conveniently test whether or not a point belongs to the risk region (or some other set satisfying the aggregation condition).

In *aggregation sampling*, the user specifies a number of scenarios to be in the risk region. The algorithm then draws samples from the distribution, storing those samples which lie in the risk region

and aggregating those in the non-risk region into a single point. In particular, the samples in the non-risk region are aggregated into their mean. The algorithm terminates when the specified number of risk scenarios has been reached. This is detailed in Algorithm 1. In *aggregation reduction* one draws a fixed number of samples from the distribution and then aggregates all those in the non-risk region.

input : $\mathcal{R} \subset \mathbb{R}^d$ set satisfying aggregation condition, $N_{\mathcal{R}}$ number of required risk scenarios
output: $\{(y_s, p_s)\}_{s=1}^{N_{\mathcal{R}}+1}$ scenario set
$n_{\mathcal{R}^c} \leftarrow 0$, $n_{\mathcal{R}} \leftarrow 0$, $y_{\mathcal{R}^c} = \mathbf{0}$;
while $n_{\mathcal{R}} < N_{\mathcal{R}}$ **do**
  Sample new point y;
  if $y \in \mathcal{R}$ **then**
    $n_{\mathcal{R}} \leftarrow n_{\mathcal{R}} + 1$;
    $y_{n_{\mathcal{R}}} \leftarrow y$;
  end
  else
    $n_{\mathcal{R}^c} \leftarrow n_{\mathcal{R}^c} + 1$;
    $y_{\mathcal{R}^c} \leftarrow \frac{1}{n_{\mathcal{R}^c}+1}\left(n_{\mathcal{R}^c} y_{\mathcal{R}^c} + y\right)$
  end
end
foreach i *in* $1, \ldots, N_{\mathcal{R}}$ **do** $p_i \leftarrow \frac{1}{(n_{\mathcal{R}^c}+N_{\mathcal{R}})}$ **if** $n_{\mathcal{R}^c} > 0$ **then**
  $y_{n_{\mathcal{R}^c}+1} \leftarrow \frac{\mathbf{s}}{n_{\mathcal{R}^c}}$;
end
else
  Sample new point y;
  $n_{\mathcal{R}^c} \leftarrow 1$;
  $y_{N_{\mathcal{R}}+1} \leftarrow y$;
end
$p_{N_{\mathcal{R}}+1} \leftarrow \frac{n_{\mathcal{R}^c}}{n_{\mathcal{R}^c}+N_{\mathcal{R}}}$

Algorithm 1: Aggregation sampling

Aggregation sampling and aggregation reduction can be thought of as equivalent to sampling from the aggregated random vector for large sample sizes. Therefore, aggregation sampling and aggregation reduction are consistent with sampling only if $\mathcal{R}$ satisfies the aggregation condition and $\mathbb{E}\left[Y|Y \in \mathcal{R}^c\right] \in \mathcal{R}^c$. For the precise conditions required for consistency and proofs of these see Theorem 4.4.

We now study the performance of our methodology. Let q the probability of the non-risk region, and n the desired number of risk scenarios. Let $N(n)$ denote the *effective sample size* for aggregation sampling, that is, the number of samples drawn until the algorithm terminates[2]. The aggregation sampling algorithm can be considered as a sequence of Bernoulli trials where a trial is a success if the corresponding sample lies in the non-risk region, and which terminates once we have reached n failures, that is, once we have sampled n scenarios from the risk region. We can therefore write down the distribution of $N(n)$ as follows:

$$N(n) \sim n + \mathcal{NB}(n, q),$$

where $\mathcal{NB}(N, q)$ denotes a *negative binomial* random variable. The expected effective sample size of aggregation sampling is thus as follows:

$$\mathbb{E}\left[N(n)\right] = n + n\frac{q}{1-q} \tag{11}$$

Aggregation reduction can be viewed as a sequence of n Bernoulli trials, where success and failure are defined in the same way as described above. The number of scenarios in the reduced sample, $R(n)$ is as follows:

$$R(n) \sim n - \mathcal{B}(n, q) + 1$$

[2]For simplicity of exposition we discount the event that the while loop of the algorithm terminates with $n_{\mathcal{R}^c} = 0$ which occurs with probability q^n

where $\mathcal{B}(n, q)$ denotes a binomial random variable. The expected reduction in scenarios in aggregation reduction is thus $nq - 1$.

For both aggregation sampling and aggregation reduction we can see that the efficacy improves as the probability of the non-risk region increases. In particular, given the general properties of risk regions (3) and (4), we can expect the performance of our methods to improve as β, the level of tail risk measure increases, and as $\mathcal{X}$, our feasible region of decisions becomes more constrained.

Because aggregation sampling and aggregating reduction only represents the non-risk region with a single scenario, it does not in general preserve the overall expectation of the cost function, or any other statistics except for the value of a β-tail risk measure. These algorithms should therefore generally only be used for problems which only involve β-tail risk measures. However, if the cost function is affine (in the sense of Corollary 2.13), then collapsing all points in the non-risk region to the conditional expectation preserves the overall expectation.

If expectation or any other statistic of the cost function is used in the optimization problem then one could represent the non-risk region region with many scenarios. For example, instead of aggregating all scenarios in the non-risk region into a single point we could apply a clustering algorithm to them such as k-means. Such a clustered scenario set for the portfolio selection problem is illustrated in Figure 2; see Section 5 for details of this problem. The ideal allocation of points between the risk and non-risk regions will be problem dependent and is beyond the scope of this paper.

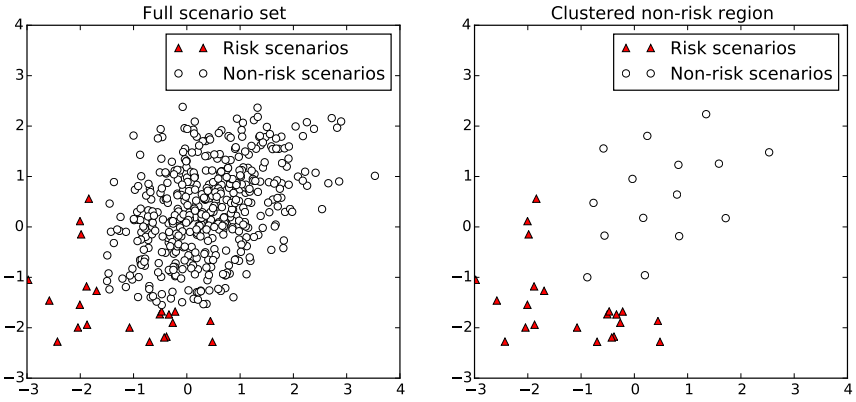

Figure 2: Scenario reduction via k-means clustering on a non-risk region for a portfolio selection problem

4 Consistency of aggregation sampling

The reason that aggregation sampling and aggregation reduction work is that for large sample sizes, they are equivalent to sampling from the aggregated random vector, and if the aggregation condition holds then the aggregated random vector yields the same optimization problem as the original random vector. We only prove consistency for aggregation sampling and not aggregation reduction as the proofs are very similar. Essentially, the only difference is that aggregation sampling has the additional complication of terminating after a random number of samples.

We suppose in this section that we have a sequence of independently identically distributed (i.i.d.) random vectors $Y_1, Y_2, \ldots$ with the same distribution as Y, and which are defined on the product probability space Ω^∞.

4.1 Uniform convergence of empirical β-quantiles

The i.i.d. sequence of random vectors $Y_1, Y_2, \ldots$ can be used to estimate the distribution and quantile functions of Y. We introduce the additional short-hand for the empirical distribution and quantile

functions:

$$F_{n,x}(z) := \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}_{\{f(x,Y_i)\le z\}},$$
$$F^{-1}_{n,x}(\beta) := \inf\{z \in \mathbb{R} :\ F_{n,x}(z) \ge \beta\}.$$

Note that these are random-valued functions on the probability space Ω^∞. It is immediate from the strong law of large numbers that for all $\bar{x} \in \mathbb{R}$ and $z \in \mathbb{R}$, we have $F_{n,x}(z) \overset{\text{w.p.1}}{\to} F_{\bar{x}}(z)$ as $n \to \infty$. In addition, if $F_{\bar{x}}$ is strictly increasing at $z = F^{-1}_{\bar{x}}$ then we also have $F^{-1}_{n,\bar{x}}(\beta) \overset{\text{w.p.1}}{\to} F^{-1}_{\bar{x}}(\beta)$ as $n \to \infty$; see for instance [Ser80][Chapter 2]. The following result extends this pointwise convergence to a convergence result which is uniform with respect to $x \in \mathcal{X}$.

Theorem 4.1. Suppose the following hold:

(i) For each $x \in \mathcal{X}$, F_x is strictly increasing and continuous in some neighborhood of $F^{-1}_x(\beta)$

(ii) For all $\bar{x} \in \mathcal{X}$ the mapping $x \mapsto f(x, Y)$ is continuous at $\bar{x}$ with probability 1.

(iii) $\mathcal{X} \subset \mathbb{R}^k$ is compact

then $F^{-1}_{n,x}(\beta) \to F^{-1}_x(\beta)$ uniformly on $\mathcal{X}$ with probability 1.

The proof of this result relies on various continuity properties of the distribution and quantile functions which are are provided in Appendix A. Some elements of the proof below have been adapted from [SDR09, Theorem 7.48], a result which concerns the uniform convergence of expectation functions.

Proof. Fix $\epsilon_0 > 0$ and $\bar{x} \in \mathcal{X}$. Since $F_{\bar{x}}$ is continuous in a neighborhood of $F^{-1}_{\bar{x}}(\beta)$, there exists $0 < \epsilon < \epsilon_0$ such $F_{\bar{x}}$ is continuous at $F^{-1}_{\bar{x}}(\beta) \pm \epsilon$. Since $F_{\bar{x}}$ is strictly increasing at $F^{-1}_{\bar{x}}(\beta)$,

$$\delta := \min\{\beta - F_{\bar{x}}\left(F^{-1}_{\bar{x}}(\beta) - \epsilon\right),\ F_{\bar{x}}\left(F^{-1}_{\bar{x}}(\beta) + \epsilon\right) - \beta\} > 0.$$

By Corollary A.2 the mapping $x \mapsto F_x\left(F^{-1}_{\bar{x}}(\beta) - \epsilon\right)$ is continuous at $\bar{x}$ with probability 1. Applying Lemma A.4, there exists a neighborhood W of $\bar{x}$ such that with probability 1, for n large enough

$$\sup_{x \in W \cap \mathcal{X}} \left|F_{n,x}(F^{-1}_{\bar{x}}(\beta) - \epsilon) - F_{n,\bar{x}}(F^{-1}_{\bar{x}}(\beta) - \epsilon)\ \right| < \frac{\delta}{2}.$$

In addition, by the strong law of large numbers, with probability 1, for n large enough

$$\left|F_{n,\bar{x}}\left(F^{-1}_{\bar{x}}(\beta) - \epsilon\right) - F_{\bar{x}}\left(F^{-1}_{\bar{x}}(\beta) - \epsilon\right)\right| < \frac{\delta}{2} \tag{12}$$

Thus, for all $x \in W \cap \mathcal{X}$ we have that

$$\left|F_{n,x}\left(F^{-1}_{\bar{x}}(\beta) - \epsilon\right) - F_{\bar{x}}\left(F^{-1}_{\bar{x}}(\beta) - \epsilon\right)\right| < \delta.$$

Similarly, we we can choose W so that we also have

$$\left|F_{n,x}\left(F^{-1}_{\bar{x}}(\beta) + \epsilon\right) - F_{\bar{x}}\left(F^{-1}_{\bar{x}}(\beta) + \epsilon\right)\right| < \delta.$$

and so

$$F_{n,x}\left(F^{-1}_{\bar{x}}(\beta) - \epsilon\right) < \beta < F_{n,x}\left(F^{-1}_{\bar{x}}(\beta) + \epsilon\right).$$

Hence, we have that with probability 1, for n large enough

$$\sup_{x \in W \cap \mathcal{X}} \left|F^{-1}_{n,x}(\beta) - F^{-1}_{\bar{x}}(\beta)\right| \le \epsilon < \epsilon_0. \tag{13}$$

Also, by Proposition A.3 the function $x \mapsto F^{-1}_x(\beta)$ is continuous and so the neighborhood can also be chosen so that

$$\sup_{x \in W \cap \mathcal{X}} \left|F^{-1}_{\bar{x}}(\beta) - F^{-1}_x(\beta)\right| < \epsilon_0, \tag{14}$$

and so combining (13) and (14) we have

$$\sup_{x\in W\cap\mathcal{X}} \left|F_{n,x}^{-1}(\beta) - F_x^{-1}(\beta)\right| < 2\epsilon_0.$$

Finally, since $\mathcal{X}$ is compact, there exists a finite number of points $x_1, \ldots, x_m \in \mathcal{X}$ with corresponding neighborhoods $W_1, \ldots, W_m$ covering $\mathcal{X}$, such that with probability 1, for n large enough the following holds:

$$\sup_{x\in W_j\cap\mathcal{X}} \left|F_{n,x}^{-1}(\beta) - F_x^{-1}(\beta)\right| < 2\epsilon_0 \qquad \text{for } i = 1, \ldots, m$$

that is, with probability 1, for n large enough

$$\sup_{x\in\mathcal{X}} \left|F_{n,x}^{-1}(\beta) - F_x^{-1}(\beta)\right| < 2\epsilon_0.$$

□

In the next subsection this result will be used to show that any point in the interior of the non-risk region will, with probability 1, be in the non-risk region of the sampled scenario set for a large enough sample size.

4.2 Equivalence of aggregation sampling with sampling from aggregated random vector

The main obstacle in showing that aggregation sampling is equivalent to sampling from the aggregated random vector is to show that the aggregated scenario in the non-risk region converges almost surely to the conditional expectation of the non-risk region as the number of specified risk scenarios tends to infinity. Recall from Section 3 that $N(n)$ denotes the effective sample size in aggregation sampling when we require n risk scenarios and is distributed as $n + \mathcal{NB}(n, q)$ where q is the probability of the non-risk region. The purpose of the next Lemma is to show that as $n \to \infty$ the number of samples drawn from the non-risk region almost surely tends to infinity.

Lemma 4.2. Suppose $M(n) \sim \mathcal{NB}(n, p)$ where $0 < p < 1$. Then with probability 1 we have that $\lim_{n\to\infty} M(n) = \infty$.

Proof. First note that,

$$\begin{aligned}\{\lim_{n\to\infty} M(n) = \infty\}^c &= \bigcup_{k\in\mathbb{N}} \left(\bigcap_{n\in\mathbb{N}} \bigcup_{t>n} \{M(t) > k\}^c \right) \\ &= \bigcup_{k\in\mathbb{N}} \limsup_{n\to\infty} \{M(n) \le k\}.\end{aligned}$$

Hence, to show that $\mathbb{P}(\{\lim_{n\to\infty} M(n) = \infty\}) = 1$ it is enough to show for each $k \in \mathbb{N}$ we have that

$$\mathbb{P}\left(\limsup_{n\to\infty} \{M(n) \le k\} \right) = 0. \tag{15}$$

Now, fix $k \in \mathbb{N}$. Then for all $n \in \mathbb{N}$ we have that

$$\mathbb{P}(M(n) = k) = \binom{k+n-1}{k} (1-p)^n \, p^k,$$

and in particular,

$$\begin{aligned}\mathbb{P}(M(n+1) = k) &= \binom{k+n}{k} (1-p)^{n+1} p^k \\ &= \frac{k+n}{n}(1-p)\, \mathbb{P}(M(n) = k).\end{aligned}$$

For large enough n we have that $\frac{k+n}{n}(1-p) < 1$, hence $\sum_{n=1}^{\infty} \mathbb{P}(M(n) = k) < +\infty$ and so

$$\begin{aligned}\sum_{n=1}^{\infty} \mathbb{P}(M(n) \leq k) &= \sum_{n=1}^{\infty} \sum_{j=1}^{k} \mathbb{P}(M(n) = j) \\ &= \sum_{j=1}^{k} \sum_{n=1}^{\infty} \mathbb{P}(M(n) = j) \\ &< \infty.\end{aligned}$$

The result (15) now holds by the first Borel-Cantelli Lemma [Bil95, Section 4]. □

The next Corollary shows that the strong law of large numbers still applies for the conditional expectation of the non-risk region in aggregation sampling despite the sample size being a random quantity.

Corollary 4.3. Suppose $\mathbb{E}[|Y|] < +\infty$ and $\mathbb{P}(Y \in \mathcal{R}^c) > 0$, then

$$\frac{1}{N(n) - n} \sum_{i \in 1 \ldots, N(n):\ Y_i \in \mathcal{R}^c} Y_i \to \mathbb{E}[\, Y | Y \in \mathcal{R}^c\,] \text{ with probability 1 as } n \to \infty$$

Proof. Define the following measurable subsets of Ω^{∞}:

$$\begin{aligned}\Omega_1 &= \{\omega \in \Omega : \lim_{n\to\infty} N(n)(\omega) - n = \infty\}, \\ \Omega_2 &= \{\omega \in \Omega : \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}_{\{Y_i(\omega) \in \mathcal{R}^c\}} Y_i(\omega) = \mathbb{E}\left[\mathbb{1}_{\{Y \in \mathcal{R}^c\}} Y\right]\}, \\ \Omega_3 &= \{\omega \in \Omega : \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}_{\{Y_i(\omega) \in \mathcal{R}^c\}} = \mathbb{P}(Y \in \mathcal{R}^c)\}.\end{aligned}$$

By the strong law of large numbers Ω_2 and Ω_3 have probability one. Since $N(n) - n \sim \mathcal{NB}(n, q)$, where $q = \mathbb{P}(Y \in \mathcal{R}^c)$, Ω_1 has probability 1 by Lemma 4.2. Therefore, $\Omega_1 \cap \Omega_2 \cap \Omega_3$ has probability 1 and so it is enough to show that for any $\omega \in \Omega_1 \cap \Omega_2 \cap \Omega_3$ we have that

$$\frac{1}{N(n)(\omega) - n} \sum_{i \in 1 \ldots, N(n):\ Y_i(\omega) \in \mathcal{R}^c} Y_i(\omega) \to \mathbb{E}[\, Y | Y \in \mathcal{R}^c\,] \text{ as } n \to \infty.$$

Let $\omega \in \Omega_1 \cap \Omega_2 \cap \Omega_3$. Since $\omega \in \Omega_2 \cap \Omega_3$, we have that as $m \to \infty$:

$$\begin{aligned}\frac{1}{\frac{1}{m}\sum_{i=1}^{m} \mathbb{1}_{\{Y_i(\omega) \in \mathcal{R}^c\}}} \frac{1}{m} \sum_{i=1}^{m} \mathbb{1}_{\{Y_i(\omega) \in \mathcal{R}^c\}} Y_i &\to \frac{1}{\mathbb{P}(Y \in \mathcal{R}^c)} \mathbb{E}\left[\mathbb{1}_{\{Y \in \mathcal{R}^c\}} Y\right] \\ &= \mathbb{E}[Y | Y \in \mathcal{R}^c].\end{aligned}$$

Now, fix $\epsilon > 0$. Then there exists $N_1(\omega) \in \mathbb{N}$ such

$$m > N_1(\omega) \implies \left| \frac{1}{\frac{1}{m}\sum_{i=1}^{m} \mathbb{1}_{\{Y_i(\omega) \in \mathcal{R}^c\}}} \frac{1}{m} \sum_{i=1}^{m} \mathbb{1}_{\{Y_i(\omega) \in \mathcal{R}^c\}} Y_i - \mathbb{E}[Y | Y \in \mathcal{R}^c] \right| < \epsilon.$$

Since $\omega \in \Omega_1$ there exists $N_2(\omega)$ such that

$$n > N_2(\omega) \implies N(n)(\omega) > N_1(\omega).$$

Noting that

$$\frac{1}{\frac{1}{N(n)(\omega)}\sum_{i=1}^{N(n)(\omega)}\mathbb{1}_{\{Y_i(\omega)\in\mathcal{R}^c\}}}\frac{1}{N(n)(\omega)}\sum_{i=1}^{N(n)(\omega)}\mathbb{1}_{\{Y_i(\omega)\in\mathcal{R}^c\}}Y_i(\omega)=\frac{1}{\frac{N(n)(\omega)-n}{N(n)(\omega)}}\frac{1}{N(n)(\omega)}\sum_{i=1}^{N(n)(\omega)}\mathbb{1}_{\{Y_i(\omega)\in\mathcal{R}^c\}}Y_i(\omega)$$
$$=\frac{1}{N(n)(\omega)-n}\sum_{i:Y_i(\omega)\in\mathcal{R}^c}Y_i,$$

we have that

$$n>N_2\implies\left|\frac{1}{N(n)(\omega)-n}\sum_{i:Y_i(\omega)\in\mathcal{R}^c}Y_i(\omega)-\mathbb{E}\left[Y|Y\in\mathcal{R}^c\right]\right|<\epsilon$$

and so $\frac{1}{N(n)(\omega)-n}\sum_{i:Y_i(\omega)\in\mathcal{R}^c}Y_i(\omega)\to\mathbb{E}\left[Y|Y\in\mathcal{R}^c\right]$ as $n\to\infty$. □

To show that aggregation sampling yields solutions consistent with the underlying random vector Y, we show that with probability 1, for n large enough, it is equivalent to sampling from the aggregated random vector $\psi_{\mathcal{R}}(Y)$, as defined in Definition 2.12. As long as the region $\mathcal{R}$ satisfies the aggregation condition, Theorem 2.8 tells us that $\rho_\beta\left(f(x,\psi_{\mathcal{R}}(Y))\right)=\rho_\beta\left(f(x,Y)\right)$ for all $x\in\mathcal{X}$ and so if sampling is consistent for the risk measure ρ_β, aggregation sampling also consistent.

Denote by $\tilde{F}_{n,x},\tilde{F}^{-1}_{n,x}$, the empirical distribution, and quantile functions respectively and by $\tilde{\rho}_{n,\beta}(x)$ the value of the tail-risk measure for the decision $x\in\mathcal{X}$ for the sample from the aggregated random vector: $\psi_{\mathcal{R}}(Y_1),\ldots,\psi_{\mathcal{R}}(Y_n)$. Similarly, denote by $\hat{F}_{n,x},\hat{F}^{-1}_{n,x}$, and $\hat{\rho}_{n,\beta}$ the analogous functions for the scenario set constructed by aggregation sampling with n risk scenarios. Note that these latter functions will depend on the sample $Y_1,\ldots,Y_{N(n)}$. Note also that like $F_{n,x}$ and $F^{-1}_{n,x}$, all these functions are random and defined on the same sample space Ω^∞.

Theorem 4.4. Suppose the following conditions hold:

(i) $(x,y)\mapsto f(x,y)$ is continuous on $\mathcal{X}\times\mathbb{R}^d$

(ii) For each $x\in\mathcal{X}$, F_x is strictly increasing and continuous in some neighborhood of $F_x^{-1}(\beta)$

(iii) $\mathbb{E}\left[\,Y|Y\in\mathcal{R}^c\,\right]\in\operatorname{int}\left(\mathcal{R}^c\right)$

(iv) $\mathcal{X}$ is compact.

Then, with probability 1, for n large enough $\tilde{\rho}_{n,\beta}\equiv\hat{\rho}_{N(n),\beta}$.

Proof. Note that if

$$z>\max\left\{f\left(x,\frac{1}{N(n)(\omega)-n}\sum_{i\in1\ldots,N(n)(\omega):\ Y_i(\omega)\in\mathcal{R}^c}Y_i(\omega)\right),f\left(x,\mathbb{E}\left[\,Y|Y\in\mathcal{R}^c\,\right]\right)\right\}$$

then

$$\hat{F}_{n,x}(z)(\omega)=\frac{N(n)(\omega)-n}{N(n)(\omega)}+\frac{1}{N(n)(\omega)}|\{1\le i\le N(n)(\omega)\mid f(x,Y_i(\omega))\le z\text{ and }Y_i(\omega)\in\mathcal{R}\}|=\tilde{F}_{N(n),x}(z)(\omega).$$

So if we have

$$\hat{F}^{-1}_{n,x}(\beta)(\omega)>\max\{f\left(x,\frac{1}{N(n)(\omega)-n}\sum_{i\in1\ldots,N(n)(\omega):\ Y_i(\omega)\in\mathcal{R}^c}Y_i(\omega)\right),f\left(x,\mathbb{E}\left[\,Y|Y\in\mathcal{R}^c\,\right]\right)\}\qquad(16)$$

then this implies that $\hat{F}^{-1}_{n,x}(u)(\omega)=\tilde{F}^{-1}_{N(n),x}(u)(\omega)$ for all $u\ge\beta$, which in turn implies $\hat{\rho}_{\beta,n}(x)(\omega)=\tilde{\rho}_{\beta,N(n)}(x)(\omega)$. Hence, it is enough to show that with probability 1, for sufficiently large n, the inequality (16) holds for all $x\in\mathcal{X}$.

Since $\mathbb{E}[\,Y|Y \in \mathcal{R}^c\,] \in \operatorname{int}(\mathcal{R}^c)$ we have that

$$f(x, \mathbb{E}[Y|Y \in \mathcal{R}^c]) < F_x^{-1}(\beta) \qquad \text{for all } x \in \mathcal{X}$$

and since $\mathcal{X}$ is compact there exists $\delta > 0$ such that

$$\sup_{x \in \mathcal{X}} \left(F_x^{-1}(\beta) - f\left(x, \mathbb{E}[\,Y|Y \in \mathcal{R}^c\,]\right)\right) > \delta. \tag{17}$$

The continuity of $f(x, y)$ and again the compactness of $\mathcal{X}$ implies that there exists $\gamma > 0$ such that

$$|y - \mathbb{E}[\,Y|Y \in \mathcal{R}^c\,]| < \gamma \implies \sup_{x \in \mathcal{X}} |f(x, y) - f(x, \mathbb{E}[\,Y|Y \in \mathcal{R}^c]) \,| < \frac{\delta}{2}$$

Thus, by Corollary 4.3, with probability 1, for n large enough

$$\left| f\left(x, \frac{1}{N(n) - n} \sum_{i \in 1 \ldots, N(n):\ Y_i \in \mathcal{R}^c} Y_i\right) - f\left(x, \mathbb{E}[\,Y|Y \in \mathcal{R}^c]\right) \right| < \frac{\delta}{2} \tag{18}$$

Also, by Theorem 4.1, given $N(n) > n$, for n large enough

$$\sup_{x \in \mathcal{X}} \left|F_x^{-1}(\beta) - \tilde{F}_{N(n),x}^{-1}(\beta)\right| < \frac{\delta}{2}, \tag{19}$$

which implies for all $x \in \mathcal{X}$

$$\begin{aligned}
\tilde{F}_{N(n),x}^{-1}(\beta) - f\left(x, \frac{1}{N(n) - n} \sum_{i \in 1 \ldots, N(n):\ Y_i \in \mathcal{R}^c} Y_i\right) &\geq \left(F_x^{-1}(\beta) - \frac{\delta}{2}\right) \\
&\quad - \left(f\left(x, \mathbb{E}[\,Y|Y \in \mathcal{R}^c\,]\right) + \frac{\delta}{2}\right) \qquad \text{by (18) and (19)} \\
&= \underbrace{\left(F_x^{-1}(\beta) - f\left(x, \mathbb{E}[\,Y|Y \in \mathcal{R}^c\,]\right)\right)}_{> \delta \text{ by (17)}} - \delta \\
&> 0.
\end{aligned}$$

Similarly with probability 1 for n large enough we have $\tilde{F}_{N(n),x}^{-1}(\beta) > f\left(x, \mathbb{E}[\,Y|Y \in \mathcal{R}^c\,]\right)$ for all $x \in \mathcal{X}$. Therefore the inequality (16) holds with probability 1 for sufficiently large n as required. □

5 Risk regions for the portfolio selection problem

In this section we characterize exactly the risk region for the portfolio selection problem when the asset returns are elliptically distributed. In Section 5.1 we formulate the basic problem and, to provide some intuition, we find the risk region by brute force for an arbitrary discrete distribution. In Section 5.2 we define elliptical distributions and give the non-risk region for the unconstrained problem, and finally in Section 5.3 we characterize the non-risk region when portfolios are constrained to a a convex set.

5.1 Problem statement and brute force aggregation

In the portfolio selection problem, one aims to choose a portfolio of financial assets with uncertain returns. For $i = 1, \ldots, d$, let x_i denote the amount to invest in asset i, and Y_i the random return of asset i. The loss function in this problem is the negative total return, that is $f(x, Y) = \sum_{i=1}^{d} -x_i Y_i = -x^T Y$. The optimization problem will typically try to balance the expected profit against the risk in some way, and so our problem is usually of one of the following forms:

$$
\begin{aligned}
&\text{(i)} \underset{x\in\mathcal{X}}{\text{minimize}}\ \rho_\beta(-x^TY) \\
&\qquad\qquad \text{subject to } \mathbb{E}\left[x^TY\right] \geq t \\
&\text{(ii)} \underset{x\in\mathcal{X}}{\text{maximize}}\ \mathbb{E}\left[x^TY\right] \\
&\qquad\qquad \text{subject to } \rho_\beta(-x^TY) \leq s \\
&\text{(iii)} \underset{x\in\mathcal{X}}{\text{minimize}}\ \rho_\beta(-x^TY) + v\mathbb{E}\left[-x^TY\right]
\end{aligned}
$$

where $v \geq 0$ and $\mathcal{X} \subset \mathbb{R}^d$ represents the set of valid portfolios. The set $\mathcal{X}$ of feasible portfolios may encompass constraints like no short-selling ($x \geq 0$), total investment ($\sum_{i=1}^d x_i = 1$) and quotas on certain stocks or combinations of stocks ($x \leq c$).

For a given portfolio $x \in \mathcal{X}$, the corresponding risk region is the half-space of points where loss is greater than or equal to the β-quantile:

$$\mathcal{R}_x = \{y \in \mathbb{R}^d :\ -x^Ty \geq F_x^{-1}(\beta)\}$$

For a discrete distribution of returns, finding the β-quantile of the loss associated to a particular portfolio is a case of ordering all scenarios according to their loss and selecting the appropriate order statistic. In Figure 3 we have illustrated a scenario set of returns for two hypothetical assets. The line in this figure separates all those scenarios with loss below the β-quantile from those with loss above for the portfolio $x = (\frac{1}{2}, \frac{1}{2})$.

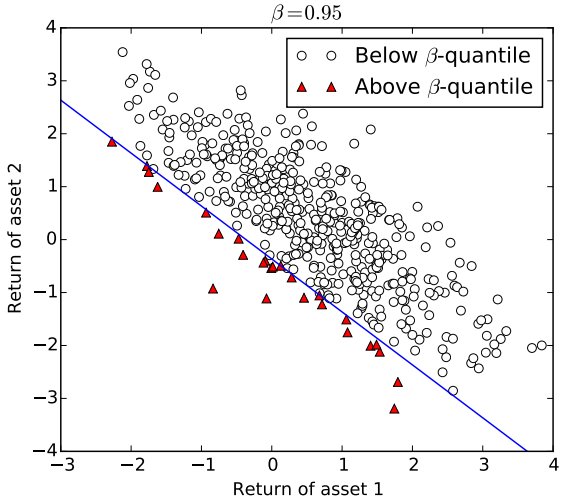

Figure 3: Scenarios with loss above and below β-quantile for one portfolio

Recall that the risk region associated to a set of feasible decisions is the union of all risk regions for decisions in that set. Thus, we can find all scenarios in the risk region by calculating the β-quantile for all feasible portfolios. On the left hand side of Figure 4, for the same scenario set in Figure 3, we have identified the risk scenarios for the set of feasible portfolios $\mathcal{X} = \{(x_1, x_2) \in \mathbb{R}^2 :\ x_1, x_2 \geq 0,\ x_1 + x_2 = 1\}$.

Corollary 2.13 states that if the aggregation condition holds, then all the mass in the non-risk region can be aggregated into its conditional expectation without affecting the value of the expectation of the loss or any tail risk measure. For a discrete distribution, we noted in Section 2.2 that the aggregation condition always holds for the risk region. On the right-hand side of Figure 4 is illustrated the same scenario set where all non-risk scenarios have been aggregated into a single point. Note that the β-quantile lines have not changed after aggregation. By aggregating all the non-risk scenarios into a single point we substantially reduce the computational cost of solving the corresponding portfolio selection problem.

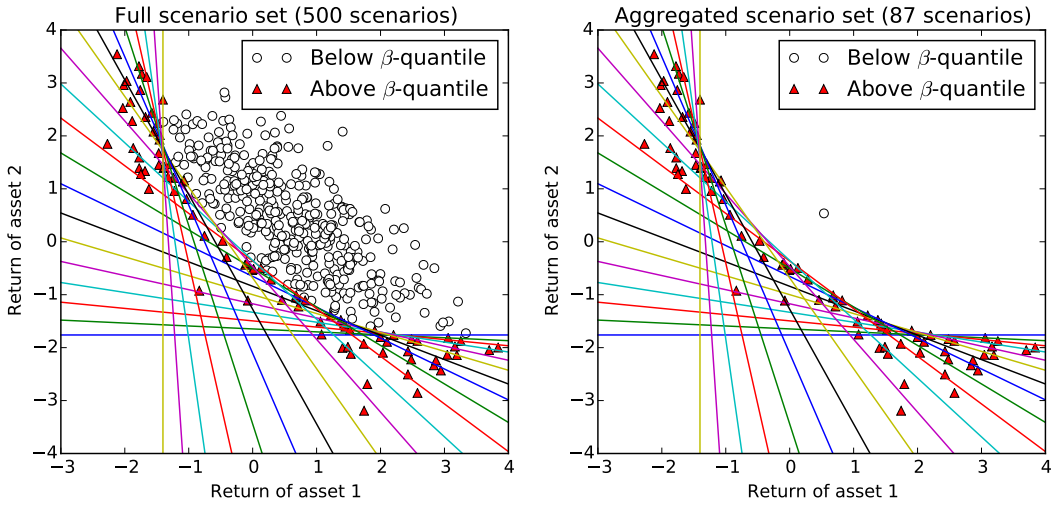

Figure 4: Scenario set separated into risk and non-risk scenarios: full scenario set (left) and aggregated scenario set (right)

The following corollary gives sufficient conditions for the risk region to satisfy the aggregation condition for continuous distributions.

Corollary 5.1. Suppose that $\mathcal{Y} = \mathbb{R}^d$ and there exist $x_1, x_2 \in \mathcal{X}$ which are linearly independent. Then, for any $\mathcal{R} \supseteq \mathcal{R}_{\mathcal{X}}$, $\mathcal{R}$ satisfies the aggregation condition. Moreover, if $\mathcal{R}$ is convex, Y is continuous and $\mathcal{X}$ is compact, then aggregation sampling with respect to $\mathcal{R}$ is consistent in the sense of Theorem 4.4.

Proof. For the first part of this result, it is enough to show that $\mathcal{R}_{\mathcal{X}}$ satisfies the aggregation condition. We prove this by showing that all the conditions of Proposition 2.10 hold. Note that $x \mapsto -x^T y$ is continuous so condition (ii) holds immediately.

For all $x \in \mathcal{X}$ the interior of the corresponding risk region and non-risk region are open half-spaces:

$$\operatorname{int}(\mathcal{R}_x) = \{y \in \mathbb{R}^d : -x^T y > F_x^{-1}(\beta)\}$$
$$\operatorname{int}(\mathcal{R}_x^c) = \{y \in \mathbb{R}^d : -x^T y < F_x^{-1}(\beta)\}$$

Fix $\bar{x} \in \mathcal{X}$. Then either $\bar{x}$ is linearly independent to x_1 or it is linearly independent to x_2. Assume it is linearly independent to x_1. Now, $\operatorname{int}(\mathcal{R}_{\bar{x}})$ and $\operatorname{int}(\mathcal{R}_{x_1})$ are non-parallel half-spaces and so both $\operatorname{int}(\mathcal{R}_{\bar{x}} \cap \mathcal{R}_{x_1})$ and $\operatorname{int}(\mathcal{R}_{x_1} \setminus \mathcal{R}_{\bar{x}}) = \operatorname{int}(\mathcal{R}_{x_1}) \cap \operatorname{int}(\mathcal{R}_{\bar{x}}^c)$ are non-empty so condition (iii) is satisfied.

Since $\mathcal{R}_{x_1}$ and $\mathcal{R}_{x_2}$ are non-parallel half-spaces, their union $\mathcal{R}_{x_1} \cup \mathcal{R}_{x_2}$ is connected. Similarly, for any $x \in \mathcal{X}$, we must have $\mathcal{R}_x$ being non-parallel with either $\mathcal{R}_{x_1}$ or $\mathcal{R}_{x_2}$ and so $\mathcal{R}_x \cup \mathcal{R}_{x_1} \cup \mathcal{R}_{x_2}$ must also be connected. Hence, $\mathcal{R}_{\mathcal{X}} = \bigcup_{x \in \mathcal{X}} (\mathcal{R}_x \cup \mathcal{R}_{x_1} \cup \mathcal{R}_{x_2})$ is connected so condition (i) is also satisfied.

It now remains to show that aggregation sampling is consistent in the sense of Theorem 4.4. Conditions (i) and (iv) of this theorem hold trivially. Condition (iii) also holds immediately since $\mathcal{R}$ is convex, so it only remains to verify condition (ii). Since Y is continuous and has support $\mathcal{Y} = \mathbb{R}^d$, Y has a density f such that $f(y) > 0$ for all $y \in \mathbb{R}^d$. Hence, for all $x \in \mathcal{X}$, the function F_x is continuous and increasing everywhere. □

In the illustrative example above we used brute force to test whether or not a point belonged to the risk region. This approach requires the calculation of the β-quantile for all feasible decisions, which is roughly equivalent to the computational cost required to enumerate the value of the tail-risk measure for all feasible decisions. To benefit from risk regions, we instead need a convenient method to test whether or not a point belongs to it.

5.2 Non-risk region for elliptically distributed returns

By exploiting the structure of a parametric distribution, it may be possible to characterize its associated risk region in a more convenient manner. In this section we do this for elliptically distributed returns.

Elliptical distributions are a general class of distributions which include among others the multivariate Normal and multivariate t-distributions. See [FKN89] for a full overview of the subject.

Definition 5.2 (Spherical and Elliptical Distributions)**.** Let X be a random vector in $\mathbb{R}^d$, then X is said to be *spherical* if its distribution is invariant under orthonormal transformations; that is, if

$$X \sim UX \qquad \text{for all } U \in \mathbb{R}^{d\times d} \text{ orthonormal.}$$

Let Y be a random vector in $\mathbb{R}^d$, then Y is said to be *elliptical* if it can be written $Y = PX + \mu$ where $P \in \mathbb{R}^{d\times d}$ is non-singular, $\mu \in \mathbb{R}^d$, and X is random vector with spherical distribution. We will denote this $Y \sim \text{Elliptical}(X, P, \mu)$.

We will assume throughout that Y is continuous and $\mathcal{Y} = \mathbb{R}^d$ so that we can apply Corollary 5.1. An important property of elliptical distributions is that for any random vector with such a distribution, we can characterize exactly the distribution of any linear combination of the components of vector. That is, for an an elliptical distribution $Y \sim \text{Elliptical}(X, P, \mu)$ in $\mathbb{R}^d$ and $x \in \mathbb{R}^d$ we have

$$x^T Y \sim \|Px\|\, X_1 + x^T\mu. \tag{20}$$

where $\|\cdot\|$ denotes the standard Euclidean norm. This property allows us to solve some portfolio selection problems for elliptical distributions where the risk measure is positive homogeneous and translation invariant via quadratic programming or interior point algorithms. Such risk measures include the β-VaR, β-CVaR and all coherent risk measures [ADEH99]. For more details, and a proof of (20) see [KBF07]. By (20) the β-quantile of the loss of a portfolio is as follows:

$$F_x^{-1}(\beta) = \|Px\|\, F_{X_1}^{-1}(\beta) - x^T\mu.$$

Therefore, using (2) the non-risk region for $Y \sim \text{Elliptical}(X, P, \mu)$, is the following:

$$\{y \in \mathbb{R}^d : -x^T y \le \|Px\|\, F_{X_1}^{-1}(\beta) - x^T\mu \qquad \forall x \in \mathcal{X}\} \tag{21}$$

If we take $\mathcal{X} = \mathbb{R}^d$, then it can be shown that the set (21) is in fact just an ellipsoid (see Proposition (B.1)):

$$\mathcal{R}^c_{\mathbb{R}^d} = \{y \in \mathbb{R}^d : \left(y^T - \mu\right)\Sigma^{-1}\left(y - \mu\right) \le F_{X_1}^{-1}(\beta)^2\}. \tag{22}$$

where $\Sigma = P^T P$. Note that by (4) the set $\mathcal{R}_{\mathcal{X}} \subset \mathcal{R}_{\mathbb{R}^d}$ and so $\mathcal{R}_{\mathbb{R}^d}$ always satisfies the aggregation condition. Unlike (21) this characterization in (22) allows us to easily test whether or not an arbitrary point is in the risk region.

As discussed in Section 3 on scenario generation, the greater the probability of the non-risk region, the greater the benefit of our methodology over regular sampling. To gauge the utility of our methodology we calculate the probability (22) for the Normal distribution. If $Y \sim \mathcal{N}(\mu, \Sigma)$ this can be calculated exactly:

$$\begin{aligned}\mathbb{P}\left((Y-\mu)^T\Sigma^{-1}(Y-\mu) \le \Phi^{-1}(\beta)^2\right) &= \mathbb{P}\left(X^T X \le \Phi^{-1}(\beta)^2\right) \qquad \text{where } Y = PX + \mu \\ &= \mathbb{P}\left(\chi_d^2 \le \Phi^{-1}(\beta)^2\right),\end{aligned}$$

where Φ is the distribution function of the standard Normal distribution. That is, the probability of the non-risk region is invariant to the mean and covariance and can be calculated from a χ_d^2 distribution function. In Figure 5 we have plotted how the probability of the non-risk region varies with the value of β and the dimension. It shows that as the dimension increases, the probability of the non-risk region converges to zero. This convergence is so quick that for even relatively small dimensions and high values

of β, the probability of the ellipsoid is tiny. This means that the potential benefit of aggregating scenarios using this region for reasonably sized problems would be negligible. However, as we show in the next subsection, by using the constraints of our problem we can significantly increase this probability.

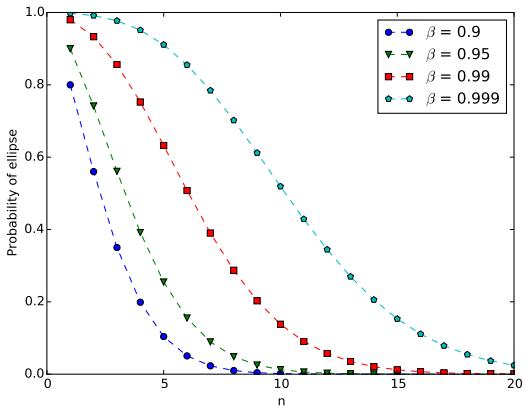

Figure 5: Plot of how mass of ellipse varies with dimension

5.3 Non-risk region with convex constraints

We now treat the more general case where the portfolios are constrained to a convex set. As well as convexity we also require the related concepts of cone and conic hull.

Definition 5.3 (Cones and Conic Hull)**.** A set $K \subset \mathbb{R}^d$ is a cone if for all $x \in K$ and $\lambda \geq 0$ we have $\lambda x \in K$. A cone is convex if for all $x_1, x_2 \in K$ and $\lambda_1, \lambda_2 \geq 0$ we have $\lambda_1 x_1 + \lambda_2 x_2 \in K$. The conic hull of a set $\mathcal{A} \subset \mathbb{R}^d$ is the smallest convex cone containing $\mathcal{A}$, and is denoted conic $(\mathcal{A})$.

The characterization of this region also makes use of the concept of a projection onto a convex set which we recall now.

Definition 5.4 (Projection)**.** Let $C \subset \mathbb{R}^d$ be a closed convex set. Then for any point $y \in \mathbb{R}^d$, we define the projection of y onto C to be the unique point $p_C(y) \in C$ such that

$$\inf_{x \in C} \|x - y\| = \|p_C(y) - y\|$$

Now, letting $K = \text{conic}\,(\mathcal{X})$, Corollary B.5 in the appendix applied to the set (21) gives us a non-risk region:

$$P^T\left(\{\tilde{y} : \|p_{K'}(\tilde{y} - \mu)\| \leq F_{X_1}^{-1}(\beta')\}\right) \tag{23}$$

where $\beta' < \beta$ and $K' = PK$. Like (22) the characterization (23) allows us easily to check whether or not a point lies in the risk region.

We now repeat our calculations of the probability of the non-risk region assuming now that $\mathcal{X} = \{x \in \mathbb{R}^d : \sum_{i=1}^d x_i = 1,\ x \geq 0\}$. The probability of the non-risk region is no longer invariant to the parameters of the Normal distribution, so for simplicity we take $\mu = 0$ and $\Sigma = I_n$. In this case we have $P = I_d$ and so $K' = K = \mathbb{R}^d_+$. Also, $p_K(y) = y_+$ where $y_+ = \max\{0, y\}$, hence

$$\mathcal{R}^c_{\mathcal{X}} = \{y \in \mathbb{R}^d : \|y_+\| \leq \Phi^{-1}(\beta)\}.$$

The probability of this region cannot be calculated analytically, and so we estimate it by Monte Carlo simulation. As Figure 6 shows, the probability of the region decays at a much slower rate as the dimension increases. This underlines the importance of making use of our constraints for finding sets which satisfy the aggregation condition.

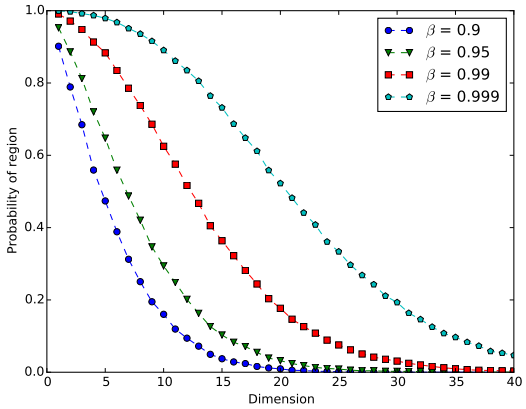

Figure 6: Plot of how probability of non-risk region varies when $K = \mathbb{R}^d_+$, and $Y \sim \mathcal{N}(0, I)$

6 Numerical tests

In this section, we test the performance of our aggregation sampling algorithm from Section 3 on the portfolio selection problem with elliptical distributions, using the non-risk region found in Section 5. We will in particular assume that the asset returns follow a Normal distribution and only impose positivity constraints on our portfolios. Note that the application of aggregation sampling in this situation is valid as all the conditions of Corollary 5.1 hold.

The test will be to compare the aggregation sampling method against basic sampling. In particular, we are interested in the quality and stability of the solutions that are yielded by our method. To this end, in each experiment we construct 50 scenario sets using sampling and aggregation sampling, solve the resulting problems, and calculate the optimality gaps for the solutions that these yield. We then estimate the probability of the non-risk region, and repeat the stability test for scenario sets of the expected effective sample size of aggregation sampling with respect to the first sample size.

We use the following problem:

$$\begin{aligned} \underset{x \geq 0}{\text{minimize}} \ & \beta\text{-CVaR}(-x^T Y) \\ \text{subject to } & \mathbb{E}\left[x^T Y\right] \geq t. \end{aligned} \tag{P}$$

where $Y \sim \mathcal{N}(\mu, \Sigma)$. In this case, it is easily shown that

$$\beta\text{-CVaR}(-x^T Y) = (1-\beta)\mu^T x + \sqrt{x^T \Sigma x} \int_{\Phi^{-1}(\beta)}^{\infty} z \, d\Phi(z) \tag{24}$$

where Φ denotes the distribution function of the standard Normal distribution. The problem (P) can be solved exactly using an interior point algorithm. We construct our Normal distributions from monthly return data between January 2007 and February 2015 from randomly selected companies in the FTSE 100 index. Note that the aggregation sampling algorithm will require the calculation of projections onto the finitely generated cones $K' = P\mathbb{R}^d_+$ where $\Sigma = P^T P$. For details on how to do this see [Ujv07].

In Figure 7 are presented the results of these stability tests for two different problems. In the first problem we have $d = 10$ and $\beta = 0.95$ and the probability of the non-risk region is estimated to be 0.594; in the second problem we have $d = 20$ and $\beta = 0.99$ for which the probability of the non-risk region is estimated to be 0.700. Note that for both of these experiments the probability of the non-risk region is much larger than for the case where asset returns are independently distributed as in Figure 6. For the first problem the expected effective sample size of aggregation sampling with 100 risk scenarios, as given by (11), is $100 + \frac{0.594}{0.406}100 \approx 246$. Similarly, the expected effective sample size of aggregation sampling for the second problem is $100 + \frac{0.7}{0.3}100 \approx 333$. In both cases, the performance of aggregation sampling for 100 risk scenarios is on a par with that of sampling for the much larger expected effective sample

size, in terms of both the quality of solutions and their stability.

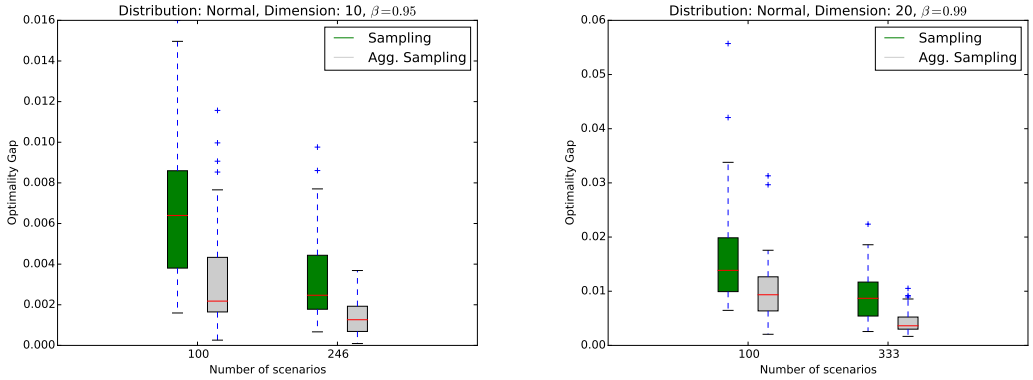

(a) Probability of risk region: 0.406

(b) Probability of risk region: 0.300

Figure 7: Optimality gap for 50 scenarios sets constructed via sampling and aggregation sampling

7 Conclusions

In this paper we have demonstrated that in stochastic programs which use a tail risk measure, a significant portion of the support of the random variables in the problem do not participate in the calculation of that tail risk measure, whatever feasible decision is used. As a consequence, for scenario-based problems, if we concentrate our scenarios in the region of the distribution which is important to the problem, the risk region, we can represent the uncertainty in our problem in a more parsimonious way, thus reducing the computational burden of solving it.

We have proposed and analyzed two specific methods of scenario generation using risk regions: aggregation sampling and aggregation reduction. Both of these methods were shown to be more effective as the probability of the non-risk region increases: in essence the higher this probability the more redundancy there is in the original distribution. Therefore, our methodology becomes more valuable as as our problem becomes more constrained, and as the level of the tail-risk increases since these changes cause the probability of the non-risk region to decrease.

However, the application of this work relies on the ability to characterize the risk region in a way which makes it convenient to test whether or not a point belongs to it. This is difficult as it depends on the cost function, the distribution of uncertain parameters, and the set of feasible decisions. An exact characterization of the risk region may not be possible for most problems, but it may be possible to find conservative regions which contain the true risk region.

For some problems the issue might be that the non-risk region has negligible probability or is even empty. Indeed we observed for the portfolio selection problem that the probability of the non-risk region quickly tended to zero as the dimension of our problem increases. A potential strategy for overcoming this problem, and more generally for improving the effectiveness of our methodology, would be the addition of artificial constraints to the problem to enlarge the non-risk region. However, even if a non-risk region has small mass, for large and difficult problems, for example those involving integer variables or with non-linear recourse problems, the reduction in computation time gained from aggregation may be significant.

In the case of the portfolio selection problem we were able to characterize the risk region in a convenient form when the distributions of asset returns are elliptical, and demonstrated the gain from aggregation sampling for simple test problems. In the paper [FTWon] we demonstrate that our methodology may be applied to more difficult and realistic portfolio selection problems such as those involving integer variables, and for which the asset returns are no longer elliptically distributed. In the same paper we also some of the technical issues involved in applying the method, such as finding the conic hull of the feasible region, and methods of projecting points onto this. We also investigate the use of artificial

constraints as a way of making our methodology more effective.

A Continuity of Distribution and Quantile Functions

Throughout we use the following set-up: $\mathcal{X} \subset \mathbb{R}^k$ a decision space, Y a random vector with support $\mathcal{Y} \subset \mathbb{R}^d$ defined on a probability space $(\Omega, \mathcal{B}, \mathbb{P})$, and a cost function $f : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$. The quantity is $f(x, Y)$ is assumed to be measurable for all $x \in \mathcal{X}$. In this appendix we prove a series of technical results related to the continuity of the distribution and quantile functions for $f(x, Y)$. These are required for the proofs in Section 4.

The following is an elementary result from the stochastic optimization literature concerning the continuity of an expectation function.

Proposition A.1. Suppose for $g : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$, and a given $\bar{x} \in \mathcal{X}$ the following holds:

(i) $x \mapsto g(x, Y)$ is continuous at $\bar{x}$ with probability 1.

(ii) There exists a neighborhood W of $\bar{x}$ and integrable $h : \mathcal{Y} \to \mathbb{R}$ such that for all $x \in W$ we have $g(x, Y) \leq h(Y)$ with probability 1.

Then, $x \mapsto \mathbb{E}\left[g(x, Y)\right]$ is continuous at $\bar{x}$.

Proof. Let $(x_k)_{k=1}^{\infty}$ be some sequence in $\mathcal{X}$ such that $x_k \to \bar{x}$ as $k \to \infty$. Without loss of generality $x_k \in W$ for all $k \in \mathbb{N}$. By assumption (i), almost surely we have $g(x_k, Y) \to f(\bar{x}, Y)$ as $k \to \infty$. Using assumption (ii) we can apply the Lebesgue theorem of dominated convergence so that:

$$\begin{aligned}\lim_{k\to\infty} \mathbb{E}\left[g(x_k, Y)\right] &= \mathbb{E}\left[\lim_{k\to\infty} g(x_k, Y)\right] \\ &= \mathbb{E}\left[g(\bar{x}, Y)\right]\end{aligned}$$

and hence $x \mapsto \mathbb{E}\left[g(x, Y)\right]$ is continuous at $\bar{x}$. □

Recall that we use the following notation for simplicity of exposition:

$$\begin{aligned}F_x(z) &:= \mathbb{P}\left(f(x, Y) \leq z\right) \\ F_x^{-1}(\beta) &:= \inf\{z \in \mathbb{R} : \ F_x(z) \geq \beta\}\end{aligned}$$

The continuity of the distribution function immediately follows from the above proposition.

Corollary A.2. Suppose for a given $\bar{x} \in \mathcal{X}$ that $x \mapsto f(x, Y)$ is continuous with probability 1 at $\bar{x}$, and for $z \in \mathbb{R}$ the distribution function $F_{\bar{x}}$ is continuous at z. Then, $x \mapsto F_x(z)$ is continuous at $\bar{x}$.

Proof. Let $g(x, Y) = \mathbb{1}_{\{f(x,Y)\leq z\}}$ so that $F_x(z) = \mathbb{E}\left[g(x, Y)\right]$. The function $g(x, Y)$ is clearly dominated by the integrable function $h(Y) = 1$. It is therefore enough to show that $x \mapsto g(x, Y)$ is almost surely continuous at $\bar{x}$ as the result will then follow from Proposition A.1.

Since $F_{\bar{x}}$ is continuous at z, we must have $\mathbb{P}\left(f(\bar{x}, Y) = z\right) = 0$. Almost surely, we have that for $\omega \in \Omega$ that $x \mapsto f(x, Y(\omega))$ is continuous at $\bar{x}$. Let's first assume that $f(\bar{x}, Y(\omega)) > z$. In this case, there exist some neighborhood V of $\bar{x}$ such that $x \in V \Rightarrow f(x, Y(\omega)) > z$, which in turn implies $|g(x, Y) - g(\bar{x}, Y)| = 0$. Hence $x \mapsto g(x, Y(\omega))$ is continuous at $\bar{x}$. The same argument holds if $f(\bar{x}, Y(\omega)) < z$. Hence, with probability 1, $x \mapsto g(x, Y)$ is continuous at $\bar{x}$. □

Continuity of the quantile function follows from the continuity of the distribution function but requires that the distribution function is strictly increasing at the required quantile.

Proposition A.3. Suppose for some $\bar{x} \in \mathcal{X}$, and $z = F_{\bar{x}}^{-1}(\beta)$ that the conditions of Corollary A.2 hold, and in addition that $F_{\bar{x}}$ is strictly increasing at $F_{\bar{x}}^{-1}(\beta)$, that is for all $\epsilon > 0$

$$F_{\bar{x}}\left(F_{\bar{x}}^{-1}(\beta) - \epsilon\right) < \beta < F_{\bar{x}}\left(F_{\bar{x}}^{-1}(\beta) + \epsilon\right).$$

Then $x \mapsto F_x^{-1}(\beta)$ is continuous at $\bar{x}$.

Proof. Assume $x \mapsto F_x^{-1}(\beta)$ is not continuous at $\bar{x}$. This means there exists $\epsilon > 0$ such that for all neighborhoods W of $\bar{x}$

$$\text{there exists } x' \in W \text{ such that } \left|F_{\bar{x}}^{-1}(\beta) - F_{x'}^{-1}(\beta)\right| > \epsilon.$$

Now set,

$$\begin{aligned}\gamma &:= \min\{\beta - F_{\bar{x}}\left(F_{\bar{x}}^{-1}(\beta) - \epsilon\right), F_{\bar{x}}\left(F_{\bar{x}}^{-1}(\beta) + \epsilon\right) - \beta\} \\ &> 0 \qquad \text{since } F_{\bar{x}} \text{ strictly increasing at } F_{\bar{x}}^{-1}(\beta).\end{aligned}$$

By the continuity of $x \mapsto F_x\left(F_{\bar{x}}^{-1}(\beta)\right)$ at $\bar{x}$ there exists W a neighborhood of $\bar{x}$, such that:

$$x \in W \Longrightarrow \left|F_x\left(F_{\bar{x}}^{-1}(\beta)\right) - F_{\bar{x}}\left(F_{\bar{x}}^{-1}(\beta)\right)\right| < \gamma. \tag{25}$$

But for the x' identified above That is,

$$\begin{aligned}F_{x'}^{-1}(\beta) &< F_{\bar{x}}^{-1}(\beta) - \epsilon \\ \text{or } F_{x'}^{-1}(\beta) &> F_{\bar{x}}^{-1}(\beta) + \epsilon\end{aligned}$$

and so given that $F_{\bar{x}}$ is non-decreasing, and by the definition of γ we must have:

$$\left|F_{\bar{x}}\left(F_{\bar{x}}^{-1}(\beta)\right) - F_{\bar{x}}\left(F_{x'}^{-1}(\beta)\right)\right| \geq \gamma$$

which contradicts (25). □

Recall, that for a sequence of i.i.d. random vectors $Y_1, Y_2, \ldots$ with the same distribution as Y, we define the sampled distribution function as follows:

$$F_{n,x}(z) := \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}_{\{f(x,Y_i)\leq z\}}.$$

The final result concerns the continuity of the sampled distribution function.

Lemma A.4. Suppose for $g : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$, and $\bar{x} \in \mathcal{X}$ the conditions from Proposition A.1 hold. Then for all $\epsilon > 0$ there exists a neighborhood W, of $\bar{x}$, such that with probability 1

$$\limsup_{n\to\infty} \sup_{x\in W\cap\mathcal{X}} \left|\frac{1}{n}\sum_{i=1}^{n} g(x,Y_i) - \frac{1}{n}\sum_{i=1}^{n} g(\bar{x},Y_i)\right| < \epsilon$$

In particular, if $x \mapsto f(x,Y)$ is continuous at $\bar{x}$ with probability 1 and $F_{\bar{x}}$ is continuous at $z \in \mathbb{R}$ then for all $\epsilon > 0$ there exists a neighborhood W, of $\bar{x}$ such that with probability 1

$$\limsup_{n\to\infty} \sup_{x\in W\cap\mathcal{X}} |F_{n,x}(z) - F_{n,\bar{x}}(z)\,| < \epsilon. \tag{26}$$

Proof. Fix $\bar{x} \in \mathcal{X}$, and $\epsilon > 0$. Let $(\gamma_k)_{k=1}^{\infty}$ be any sequence of positive numbers converging to zero and define

$$\begin{aligned}V_k &:= \{x \in \mathcal{X} : \|x - \bar{x}\| \leq \gamma_k\}, \\ \delta_k(Y) &:= \sup_{x\in V_k} |g(x,Y) - g(\bar{x},Y)\,|.\end{aligned}$$

Note first that the quantity $\delta_k(Y)$ is Lebesgue measurable (see [SDR09, Theorem 7.37] for instance). By assumption (1) the mapping $x \mapsto g(x,Y)$ is continuous at $\bar{x}$ with probability 1, hence $\delta_k(Y) \to 0$ almost

surely as $k \to \infty$. Now, since $|g(x,Y)| \leq h(Y)$ we must have $|\delta_k(Y)| \leq 2h(Y)$, therefore, by the Lebesgue dominated convergence theorem, we have that

$$\lim_{k\to\infty} \mathbb{E}\left[\delta_k(Y)\right] = \mathbb{E}\left[\lim_{k\to\infty} \delta_k(Y)\right] = 0. \tag{27}$$

Note also that

$$\sup_{x\in V_k}\left|\frac{1}{n}\sum_{i=1}^{n} g(x,Y_i) - \frac{1}{n}\sum_{i=1}^{n} g(\bar{x},Y_i)\right| \leq \frac{1}{n}\sum_{i=1}^{n}\sup_{x\in V_k}|g(x,Y_i) - g(\bar{x},Y_i)|$$

and so

$$\sup_{x\in V_k}\left|\frac{1}{n}\sum_{i=1}^{n} g(x,Y_i) - \frac{1}{n}\sum_{i=1}^{n} g(\bar{x},Y_i)\right| \leq \frac{1}{n}\sum_{i=1}^{n}\delta_k(Y_i).$$

Since the sequence of random vectors $Y_1, Y_2, \ldots$ is i.i.d. we have by the strong law of large numbers that the right-hand side of (28) converges with probability 1 to $\mathbb{E}\left[\delta_k(Y)\right]$ as $n \to \infty$. Hence, with probability 1

$$\limsup_{n\to\infty}\sup_{x\in V_k}\left|\frac{1}{n}\sum_{i=1}^{n} g(x,Y_i) - \frac{1}{n}\sum_{i=1}^{n} g(\bar{x},Y_i)\right| \leq \mathbb{E}\left[\delta_k(Y)\right]. \tag{28}$$

By (27) we can pick $k \in \mathbb{N}$ such that $\mathbb{E}\left[\delta_k(Y)\right] < \epsilon$ and so setting $W = V_k$ we have by (28) with probability 1

$$\limsup_{n\to\infty}\sup_{x\in W\cap\mathcal{X}}\left|\frac{1}{n}\sum_{i=1}^{n} g(x,Y_i) - \frac{1}{n}\sum_{i=1}^{n} g(\bar{x},Y_i)\right| < \epsilon.$$

The result (26) follows immediately as the special case $g(x,Y) = \mathbb{1}_{\{f(x,Y)\leq z\}}$. □

B Convex cone results

The results in this appendix relate to the characterization of the non-risk region for the portfolio selection problem with elliptically distributed returns. This first result allows for an exact characterization of this region for the unconstrained portfolio selection problem.

Proposition B.1. Suppose $\alpha > 0, \mu \in \mathbb{R}^d$ and $P \in \mathbb{R}^{d\times d}$. Then, for all $y \in \mathbb{R}^d$:

$$\left(y^T - \mu\right)\Sigma^{-1}\left(y-\mu\right) \leq \alpha^2 \iff x^T\left(y-\mu\right) \leq \|Px\|\ \alpha \qquad \forall x \in \mathbb{R}^d, \tag{29}$$

where $\Sigma = P^T P$.

Proof. Assume without loss of generality that $\mu = 0$. So we have to prove:

$$y^T\Sigma^{-1}y \leq \alpha^2 \iff x^T y \leq \sqrt{x^T\Sigma x}\ \alpha\ \forall x \in \mathbb{R}^d. \tag{30}$$

We first prove the forward implication. We do this by proving the converse, that is, we suppose for some $y \in \mathbb{R}^d$ that there exists $\tilde{x} \in \mathbb{R}^d$ such that $\tilde{x}^T y > \|P\tilde{x}\|\ \alpha$. First, set $y_0 = \frac{\Sigma\tilde{x}\alpha}{\|P\tilde{x}\|} = \frac{\Sigma\tilde{x}\alpha}{\sqrt{\tilde{x}^T\Sigma\tilde{x}}}$. Now,

$$y_0\Sigma^{-1}y_0 = \frac{\tilde{x}^T\Sigma^T\Sigma^{-1}\Sigma\tilde{x}\alpha^2}{\tilde{x}^T\Sigma\tilde{x}} = \alpha^2$$

and

$$\begin{aligned}\tilde{x}^T y_0 &= \tilde{x}^T\frac{\Sigma\tilde{x}\alpha}{\sqrt{\tilde{x}^T\Sigma\tilde{x}}}\\ &= \sqrt{\tilde{x}^T\Sigma\tilde{x}}\ \alpha\\ &= \|P\tilde{x}\|\ \alpha\end{aligned}$$

That is, y_0 satisfies the inequalities of this Proposition with equality. Note that we also have,

$$(y - y_0)^T \Sigma^{-1} (y - y_0) > 0 \qquad \text{since } \Sigma^{-1} \text{ is positive definite.}$$

Expanding and rearranging this expression we have,

$$\begin{aligned}
& y^T \Sigma^{-1} y - 2 y_0^T \Sigma^{-1} y + y_0 \Sigma^{-1} \tilde{y_0} > 0 \\
\Leftrightarrow \quad & y^T \Sigma^{-1} y - 2 \frac{\alpha}{\sqrt{\tilde{x}^T \Sigma \tilde{x}}} \tilde{x}^T \Sigma \Sigma^{-1} y + \alpha^2 > 0 \\
\Leftrightarrow \quad & y^T \Sigma^{-1} y - 2 \frac{\alpha}{\sqrt{\tilde{x}^T \Sigma \tilde{x}}} \tilde{x}^T y + \alpha^2 > 0 \\
\Leftrightarrow \quad & y^T \Sigma^{-1} y - 2 \frac{\alpha}{\sqrt{\tilde{x}^T \Sigma \tilde{x}}} \tilde{x}^T y + \alpha^2 > 0 \\
\Rightarrow \quad & y \Sigma^{-1} y > \alpha^2 \qquad \text{since } \tilde{x}^T y > \ \|Px\| \ \alpha,
\end{aligned}$$

as required.

We now prove the backwards implication. We again do this by proving the converse, in this case, that if $y^T \Sigma^{-1} y > \alpha^2$ then there exists $\tilde{x} \in \mathbb{R}^d \setminus \{0\}$ such that $\tilde{x}^T y > \sqrt{x^T \Sigma x} \alpha$.

Let $\tilde{x} = \Sigma^{-1} y$. Now,

$$\begin{aligned}
\tilde{x}^T y &= y^T \Sigma^{-1} y \\
&= \underbrace{\sqrt{y^T \Sigma^{-1} y}}_{= \sqrt{\tilde{x}^T \Sigma \tilde{x}}} \underbrace{\sqrt{y^T \Sigma^{-1} y}}_{> \alpha} \\
&> \sqrt{\tilde{x}^T \Sigma \tilde{x}} \ \alpha \\
&= \|Px\| \ \alpha,
\end{aligned}$$

as required. □

The following two propositions give properties about projections onto convex cones which are required in the proof of the main results of this appendix.

Proposition B.2. Suppose $K \subset \mathbb{R}^d$ is a convex cone, then, for all $y \in \mathbb{R}^d$:

$$p_K(y)^T \left(y - p_K(y) \right) = 0$$

Proof. First note that we must have $p_K(y)^T y \geq 0$. If this is not the the case then

$$\begin{aligned}
\|y - p_K(y)\|^2 &= \|p_K(y)\|^2 - 2 p_K(y)^T y + \|y\|^2 \\
&> \|y\|^2 = \|y - 0\|^2
\end{aligned}$$

which contradicts the definition of $p_K(y)$ since $0 \in K$. Now assume that $p_K(y)^T \left(y - p_K(y) \right) \neq 0$, and set $\tilde{x} = \frac{p_K(y)^T y}{\|p_K(y)\|^2} p_K(y) \in K$. Now,

$$\begin{aligned}
p_K(y)^T (\tilde{x} - y) &= p_K^T y - p_K^T y \\
&= 0.
\end{aligned}$$

By assumption $p_k^T y \neq \|p_K(y))\|^2$, and so $\tilde{x} \neq p_K(y)$, hence

$$\begin{aligned}
\|p_K(y) - y\|^2 &= \|(p_K(y) - \tilde{x}) + (\tilde{x} - y)\|^2 \\
&= \|(p_K(y) - \tilde{x})\|^2 - 2 \underbrace{(p_K(y) - \tilde{x})^T (\tilde{x} - y)}_{=0} + \|(\tilde{x} - y)\|^2 \\
&> \|(\tilde{x} - y)\|^2
\end{aligned}$$

which, again, contradictions the definition of $p_K(y)$ since $\tilde{x} \in K$. □

Proposition B.3. Let $K \subset \mathbb{R}^d$ be a convex cone and $x \in K$. Then for any $y \in \mathbb{R}^d$

$$x^T y \leq x^T p_K(y).$$

Proof. The result holds trivially if $y \in K$ so we assume $y \notin K$. Assume there exists $\tilde{x} \in K$ such that $\tilde{x}^T y > \tilde{x}^T p_K(y)$. For all $0 \leq \lambda \leq 1$ we have $\lambda x + (1-\lambda)p_K(y) \in K$. Now,

$$\begin{aligned}\|(\lambda\tilde{x} + (1-\lambda)p_K(y)) - y\|^2 - \|y - p_K(y)\|^2 &= \|\lambda(\tilde{x} - p_K(y)) + (p_K(y) - y)\|^2 - \|y - p_K(y)\|^2 \\ &= \lambda^2 \|\tilde{x} - p_K(y)\|^2 + 2\lambda(\tilde{x} - p_K(y))^T(p_K(y) - y) \\ &= \lambda^2 \|\tilde{x} - p_K(y)\|^2 - 2\lambda \underbrace{\tilde{x}^T(y - p_K(y))}_{>0 \text{ by assumption}}.\end{aligned}$$

That is, for $0 < \lambda < \frac{\tilde{x}^T(y - p_K(y))}{2\|p_K(y) - \tilde{x}\|}$ we have $\|\lambda\tilde{x} + (1-\lambda)p_K(y) - y\| < \|y - p_K(y)\|$ which contradicts the definition of $p_K(y)$. □

The next two results generalize Proposition B.1 to the case where $x \in \mathbb{R}^d$ is restricted to a convex cone. The first describes the region in the case where $P = I$, and the second generalizes the result to any non-singular matrix. In particular, it is Corollary B.5 that allows us to characterize the maximal non-risk region of portfolio selection problem for a convex feasible region.

Theorem B.4. Let $\mathcal{X} \subset \mathbb{R}^d$ be convex, and let

$$\mathcal{A} := \{y : x^T y \leq \|x\| \ \alpha \ \forall x \in \mathcal{X}\}$$

and

$$\mathcal{B} := \{y : \|p_K(y)\| \leq \alpha\}$$

where $K = \text{conic}\,(\mathcal{X})$. Then, $\mathcal{A} = \mathcal{B}$.

Proof. $(\mathcal{B} \subseteq \mathcal{A})$
Suppose $y \in \mathcal{B}$ and let $x \in \mathcal{X}$, then $x \in K$ and so

$$\begin{aligned}x^T y &\leq x^T p_K(y) && \text{by Proposition B.3} \\ &\leq \|x\| \ \|p_K(y)\| && \text{by the Cauchy-Schwartz inequality} \\ &\leq \|x\| \ \alpha && \text{since } y \in \mathcal{B}.\end{aligned}$$

Hence $y \in \mathcal{A}$.
$(\mathcal{A} \subseteq \mathcal{B})$
Suppose $y \notin \mathcal{B}$ and set $x = p_K(y) \in K$. Now,

$$\begin{aligned}x^T y &= p_K(y)^T y \\ &= p_K(y)^T p_K(y) + p_K(y)^T(y - p_K(y)) \\ &= p_K(y)^T p_K(y) && \text{by Proposition B.2} \\ &\geq \|x\| \ \alpha && \text{since } y \notin \mathcal{B}.\end{aligned}$$

Since $\mathcal{X}$ is convex we have $x = \lambda\bar{x}$ for some $\bar{x} \in \mathcal{X}$ and so we must also have $\bar{x}^T y > \|\bar{x}\|\alpha$, hence $y \notin \mathcal{A}$. □

The projection of a point onto a cone, used in the characterization above, is illustrated in Figure 8.

Corollary B.5. Suppose K is a convex cone, and $P \in \mathbb{R}^{d \times d}$ is a non-singular matrix. Let,

$$\mathcal{A} := \{y : x^T y \leq \|Px\| \ \alpha \ \forall x \in K\}$$

and

$$\mathcal{B} := P^T\left(\{\tilde{y} : \|p_{K'}(\tilde{y})\| \leq \alpha\}\right)$$

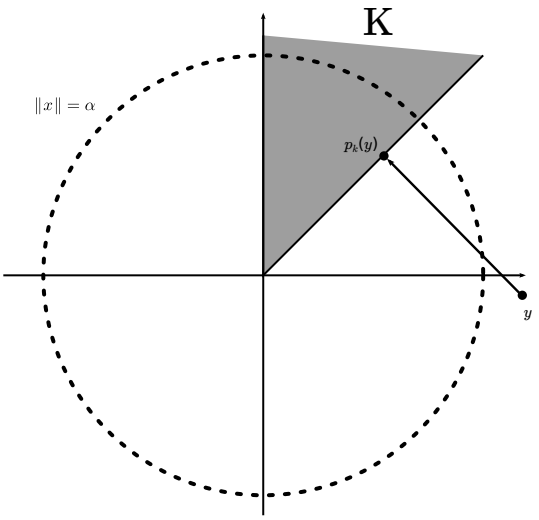

Figure 8: Projection onto a convex cone

where $K' = PK$. Then, $\mathcal{A} = \mathcal{B}$.

Proof.

$$\begin{aligned}
\mathcal{B} &= P^T\left(\{\tilde{y} : \|p_{K'}(\tilde{y})\| \leq \alpha\}\right) \\
&= P^T\left(\{\tilde{y} : \tilde{x}^T\tilde{y} \leq \|\tilde{x}\| \ \alpha \ \forall \tilde{x} \in K'\}\right) \qquad \text{by Theorem B.4} \\
&= \{y : \tilde{x}^T(P^T)^{-1}y \leq \sqrt{\tilde{x}^T\tilde{x}}\alpha \ \forall \tilde{x} \in K'\} \\
&= \{y : x^T P^T(P^T)^{-1}y \leq \|Px\| \ \alpha \ \forall x \in K\} \\
&= \{y : x^T y \leq \|Px\| \ \alpha \ \forall x \in K\} \\
&= \mathcal{A}
\end{aligned}$$

□

References

[ADEH99] P. Artzner, F. Delbaen, J. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999.

[AT02] Carlo Acerbi and Dirk Tasche. On the coherence of expected shortfall. *Journal of Banking & Finance*, 26(7):1487–1503, 2002.

[Bil95] Patrick Billingsley. *Probability and Measure*. Wiley, New York, NY, 3rd edition, 1995.

[BL97] John R. Birge and François Louveaux. *Introduction to stochastic programming*. Springer-Verlag, New York, 1997.

[FKN89] Kai-Tai Fang, Samuel Kotz, and Kai Wang Ng. *Symmetric Multivariate and Related Distributions (Chapman & Hall/CRC Monographs on Statistics & Applied Probability)*. Chapman and Hall/CRC, 11 1989.

[FTWon] Jamie Fairbrother, Amanda Turner, and Stein W. Wallace. Scenario generation for portfolio selection problems with tail risk measure, in preparation.

[GBM12] Raquel García-Bertrand and Roberto Mínguez. Iterative scenario based reduction technique for stochastic optimization using conditional value-at-risk. *Optimization and Engineering*, Online First, 2012.

[HKW03] Kjetil Høyland, Michal Kaut, and Stein W. Wallace. A heuristic for moment-matching scenario generation. *Computational Optimization and Applications*, 24(2–3):169–185, 2003.

[HR09] Holger Heitsch and Werner Römisch. Scenario tree reduction for multistage stochastic programs. *Computational Management Science*, 6(2):117–133, 2009.

[HW01] K. Høyland and S. W. Wallace. Generating scenario trees for multistage decision problems. *Management Science*, 47(2):295–307, 2001.

[Jor96] P. Jorion. *Value at Risk: The New Benchmark for Controlling Market Risk.* Irwin Professional, 1996.

[KBF07] Bahar Kaynar, Ş Ilker Birbil, and J.B.G. Frenk. Application of a general risk management model to portfolio optimization problems with elliptical distributed returns for risk neutral and risk averse decision makers. Technical report, Erasmus Research Institute of Management, 2007.

[KR93] Alan J. King and R. Tyrrell Rockafellar. Asymptotic theory for solutions in statistical estimation and stochastic programming. *Math. Oper. Res.*, 18(1):148–162, 1993.

[KW94] Peter Kall and Stein W. Wallace. *Stochastic Programming.* John Wiley & Sons, Chichester, 1994.

[KW07] Michal Kaut and Stein W. Wallace. Evaluation of scenario-generation methods for stochastic programming. *Pacific Journal of Optimization*, 3(2):257–271, 2007.

[KW11] Michal Kaut and Stein W. Wallace. Shape-based scenario generation using copulas. *Computational Management Science*, 8(1–2):181–199, 2011.

[KW12] Alan J. King and Stein W. Wallace. *Modeling with Stochastic Programming.* Springer Series in Operations Research and Financial Engineering. Springer, 2012.

[LSW06] Jeff Linderoth, Alexander Shapiro, and Stephen Wright. The empirical behavior of sampling methods for stochastic programming. *Annals of Operations Research*, 142(1):215–241, 2006.

[Mar52] H.M. Markowitz. Portfolio selection. *Journal of Finance*, 7:77–91, 1952.

[MMW99] W.K. Mak, D.P. Morton, and R.K. Wood. Monte Carlo bounding techniques for determining solution quality in stochastic programs. *Operations Research Letters*, 24:47–56, 1999.

[OR02] Włodzimierz Ogryczak and Andrzej Ruszczyński. Dual stochastic dominance and related mean-risk models. *SIAM J. Optim.*, 13(1):60–78 (electronic), 2002.

[Pfl01] G. C. Pflug. Scenario tree generation for multiperiod financial optimization by optimal discretization. *Mathematical Programming*, 89(2):251–271, 2001.

[RU00] R. Tyrrell Rockafellar and Stan Uryasev. Optimization of conditional value-at-risk. *The Journal of Risk*, 2(3):21–41, 2000.

[RU02] R. Tyrrell Rockafellar and Stan Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26(7):1443–1471, 2002.

[SDR09] Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. *Lectures on Stochastic Programming: Modeling and Theory*, volume 9 of *MPS-SIAM Series on Optimization.* SIAM, Philadelphia, 2009.

[Ser80] R.J. Serfling. *Approximation Theorems of Mathematical Statistics.* Wiley Series in Probability and Statistics - Applied Probability and Statistics Section Series. Wiley, 1980.

[Sha03] Alexander Shapiro. Monte Carlo sampling methods. In A. Ruszczyński and A. Shapiro, editors, *Stochastic Programming*, volume 10 of *Handbooks in Operations Research and Management Science*, chapter 6, pages 353–425. Elsevier Science B.V., Amsterdam, 2003.

[Tas02] Dirk Tasche. Expected shortfall and beyond. *Journal of Banking & Finance*, 26(7):1519–1533, 2002.

[Ujv07] Miklós Ujvári. On the projection onto a finitely generated cone. *Acta Cybernetica*, 2007.